Building a Nuclear Weapon

1.0 Types of Nuclear Weapons

Nuclear weapons can be grouped into different classes based on the nuclear reactions that provide their destructive energy, and on the details of their design. The popular division of nuclear weapons into fission bombs and fusion bombs is not entirely satisfactory. The spectrum of weapon design is more complex than this simple classification implies. All nuclear weapons so far invented require fission to initiate the explosive release of energy. Weapons that incorporate fusion fuel can do so in various ways, with different intended effects. This section attempts to survey the basic types of bomb designs systematically. More detailed discussions of the physics and design principles of each type will be covered in more detail in later sections.

- 1.1 Terminology
- 1.2 U.S. Nuclear Test Names
- 1.3 Units of Measurement
- 1.4 Pure Fission Weapons
- 1.5 Combined Fission/Fusion Weapons
- 1.6 Cobalt Bombs

1.1 Terminology

A variety of names are used for weapons that release energy through nuclear reactions - atomic bombs (A-bombs), hydrogen bombs (H-bombs), nuclear weapons, fission bombs, fusion bombs, thermonuclear weapons (not to mention "physics package" and "device"). A few comments about terminology is probably in order.

The earliest name for such a weapon appears to be "atomic bomb". This has been criticized as a misnomer since all chemical explosives generate energy from reactions between atoms - that is, between intact atoms consisting of both the atomic nucleus and electron shells. Further the fission weapon to which "atomic bomb" is applied is no more "atomic" than fusion weapons are. However the name is firmly attached to the pure fission weapon, and well accepted by historians, the public, and by the scientists who created the first nuclear weapons.

Since the distinguishing feature of both fission and fusion weapons is that they release energy from transformations of the atomic nucleus, the best general term for all types of these explosive devices is "nuclear weapon" (hence the name of this FAQ).

Fusion weapons are called "hydrogen bombs" (H-bombs) because isotopes of hydrogen are principal components of the nuclear reactions involved. In fact, in the earliest fusion bomb designs deuterium (hydrogen-2) was the sole fusion fuel. Fusion weapons are called "thermonuclear weapons" because high temperatures are required for the fusion reactions to occur.

1.2 Nuclear Test Names

Before discussing U.S. nuclear tests, the designation system used to identify the tests and each bomb that is tested should be clarified. Each test bomb has a code name that identifies it, the actual test has another code name. Thus the first atomic bomb was called Gadget, and it was tested in operation Trinity.

The early test operations were conducted as part of a test series, a large scale operation where many scientists, support technicians, military personnel, etc. assemble in order to set off and observe a number of devices over several weeks or months. This test series has another code name. For example the second and third test explosions of nuclear weapons (which were the fourth and fifth nuclear explosions of course) were part of the Crossroads test series. The two tests were designated Able and Baker. Sometimes the U.S. conducted two sequences of tests for different purposes jointly as a single series. When this occurred the name for each sequence was combined to form the name of the entire series (e.g. Tumbler-Snapper).

In the early test series, the same test names were reused several times. Thus there was an Able test in the Crossroads, Ranger, Buster-Jangle, and Tumbler-Snapper test series. To unambiguously identify each test the convention is to list the series code name, followed by the test name: Crossroads Able, Ranger Able, and so on. After mid-1952 unique test names began to be used, so that this convention was no longer strictly necessary. However it is useful to specify the series as well, so I have adopted the general practice of identifying tests by the series-test combination.

After 1961 the test series system was dropped as underground testing in Nevada became routine, all of which are usually considered part of the Nevada Series. Actually these tests are also still designated as being part of specific test series, but now each "series" simply corresponds to a government fiscal year (Operation Niblick is FY64, Operation Whetstone is FY65, etc.) and loses any real meaning. There was a final series of open air testing in the Pacific (the Pacific Series: Dominic I and Dominic II)in 1962, and a few special test programs (Plowshare, Vela Uniform Seismic Detonation). For Nevada Series tests after 1961, and all U.S. tests after 1963, I often follow the common practice of simply identifying them by their test names.

The code names of the actual devices are generally not well known. Many remained classified until recently (or still are). Since a bomb can only be tested once, identifying the device by the test in which it was detonated is unambiguous. In the open literature the test name has usually been used to designate the bomb that was tested, a convention followed here as well.

British tests follow a similar nomenclature but are not as systematic. Except for the first test (Hurricane), each test has both a series and test identifier. Sometimes the test identifier alone is unique, sometimes not. A variety of test series qualifiers may be attached to the series name (different from the actual test name), but the pattern for doing so is unclear. For example the Grapple series included the test Grapple 1/Short Granite (with Grapple 1 by itself uniquely identifying the test), but multiple tests all with the series designation Grapple Z. As with U.S. tests, both full series and test name designations will be given.

1.3 Units of Measurement

As a general rule I try to use the metric system throughout this FAQ. Units belonging to the U.S. Customary measurement system will crop up periodically since the primary sources of information about nuclear weapons in general are from the United States and I have not tried to convert all measurements from U.S. sources to metric yet. This presents certain problems when approximate or rounded numbers are given - should "50 lb" be rendered "22.7 kg" (falsely implying 3 digit precision) or should it be "25 kg" (and changing the measurement by 13%)?

There is also the problem of which "metric" system to use. The metric system was invented in the Napoleonic era and an international standards body has existed since 1875. Over the last two centuries many metric units had come into use, with many variations by nationality and discipline. In 1960 the SI system (Le Systeme International d'Unites) was adopted - a thoroughgoing effort to clean-up and regularize the mess that had accumulated. Many of the units that came into use early in the atomic age - curies, rems, rads, roentgens, barns, fermis, etc. - were eliminated. These units are outside the SI system, but are recognized by the SI standards as anciliary metric units to be discouraged and eventually phased out. However the original historical and scientific documents regarding the atomic age, and most writings about it since, are completely dominated by these non-SI units. Accordingly I have made no attempt to keep to the SI system.

A continuing source of confusion with nuclear weapons is the meaning of the word "ton". Normally this is used as a unit of mass or weight (a distinction I am going to ignore) in either the Metric, British Imperial, or U.S. Customary measurement systems (the last two both having two types of tons, short and long). In connection with nuclear weapons the term "ton" and its metric extensions (kiloton, megaton, etc.) are also used as units of explosive energy output or yield. The confusion is further heightened by the non-standard convention sometimes employed in the U.S. or Britain of using the abbreviation MT (or Mt, or mt) to distinguish metric tons from short tons, while also using MT (or Mt, or mt) to mean "megaton".

The SI system does not use the ton, but it does recognize a metric unit of 1000 kg called the "tonne" (aka the metric ton). In this FAQ I use tonne to mean metric ton. Ton is used either to refer to the U.S./Imperial short ton, or for the energy yield of small explosions. Which is which should be apparent from context.

Now the units of explosive energy (megatons, kilotons, or even just tons, depending on yield) are derived from attempts to compare the explosive force of a bomb to conventional explosives, the original intention was to equate it with tons of trinitrotoluene (TNT) - a workhorse military explosive. This presented problems very quickly. To which "tons" does the comparison refer? And the explosive force of TNT is not exactly a universal constant. The energy release is affected by such things as charge density, degree of confinement, temperature, and the reference end state of the explosion products. Energy outputs ranging over 980-1100 calories/g are reported.

To clarify the situation kilotons (megatons, etc.) were redefined to be a metric unit equal to exactly 10^{12} calories (4.186×10^{12} joules). Thus treating kilotons as a metric mass measurement (kilotonnes) of TNT gives a value of 1000 c/g, well within the reported range, while treating it as "kilo short tons of TNT" gives 1102 c/g, at the extreme upper end of the reported range. Thus a kiloton can be called a "kilo metric ton of TNT" and a "kilo short ton of TNT" with about equal validity.

Note that the metric definition of kiloton refers to ALL of the energy immediately released by the device, regardless of form. Although chemical explosives release essentially all of their energy as kinetic or blast energy, only part of the energy in a nuclear explosion is released in this form (though under most conditions, it is the major part). Thus a kiloton nuclear explosion actually has significantly less blast effect than a kiloton chemical explosion.

The proper capitalization of the abbreviations for kiloton and megaton is also non-standard. kt, kt, kT, and KT can all be found in the literature (the authoritative *Effects of Nuclear Weapons* avoids the issue by never abbreviating these terms). The SI standards guide recognizes the use of a lower case "t" for tonne, but no capitalization rule has ever been adopted for the explosive ton. The official SI capitalization rule for kilo is "k" and for mega it is "M" (I think it would have been more logical to reserve lower case letters for metric scale fractional units - deci, centi, milli, etc., while reserving upper case for multiple units - deka, kilo, mega, etc., but SI doesn't do this).

For no particular reason I use the lower case "t" for the explosive yield measurement ton, and I follow the SI standard in metric scale abbreviations thus we have kt and Mt for kiloton and megaton.

1.4 Pure Fission Weapons

These are weapons that only use fission reactions as a source of energy. Fission bombs operate by rapidly assembling a subcritical configuration of fissile material (plutonium or

enriched uranium) into one that is highly supercritical. The original atomic bombs tested in 16 July 1945 (device name: Gadget, test name: Trinity) and dropped on Japan in 6 August 1945 (Little Boy, over Hiroshima) and 9 August 1945 (Fat Man, over Nagasaki) were pure fission weapons.

These are the easiest nuclear weapons to design and manufacture, and the capability to do so is a prerequisite for developing any of the other weapon types. In addition to the five declared nuclear powers (the U.S., the USSR/Russia, Britain, France, and China) which have all acquired and tested these weapons, they have also been acquired by Israel, India, South Africa, and Pakistan. India has tested a fission bomb, while Israel and South Africa are suspected of having tested one.

There are practical limits to the size of pure fission bombs. Larger bombs require more fissionable material, which:

1. becomes increasingly difficult to maintain as a subcritical mass before detonation and
2. makes it harder to assemble into a high efficiency supercritical mass before stray neutrons cause predetonation.

Due to secrecy, and the boosting issue described below, it is somewhat difficult to identify the largest pure fission bomb ever tested for certain. It appears to have been the 500 kt Ivy King test by the U.S. (15 November 1952). The device exploded in this test was the Mk-18 Super Oralloy Bomb ("SOB") designed by a team led by Ted Taylor.

1.5 Combined Fission/Fusion Weapons

All nuclear weapons that are not pure fission weapons use fusion reactions to enhance their destructive effects. All weapons that use fusion require a fission bomb to provide the energy to initiate the fusion reactions. This does not necessarily mean that fusion generates a significant amount of the explosive energy, or that explosive force is even the desired effect.

1.5.1 Boosted Fission Weapons

The earliest application of fusion to useful weapons was the development of boosted fission weapons. In these weapons a few grams of a deuterium/tritium gas mixture are included in the center of the fissile core. When the bomb core undergoes enough fission, it becomes hot enough to ignite the D-T fusion reaction which proceeds swiftly. This reaction produces an intense burst of high-energy neutrons that causes a correspondingly intense burst of fissions in the core. This greatly accelerates the fission rate in the core, thus allowing a much higher percentage of the material in the core to fission before it blows apart. Typically no more than about 20% of the material in an average size pure fission bomb will split before the reaction ends (it can be much lower, the Hiroshima bomb was 1.4% efficient). By accelerating the fission process a boosted fission bomb increase the yield 100% (an unboosted 20 kt bomb can thus

become a 40 kt bomb). The actual amount of energy released by the fusion reaction is negligible, about 1% of the bomb's yield, making boosted bomb tests difficult to distinguish from pure fission tests (detecting traces of tritium is about the only way).

The first boosted weapon test was Greenhouse Item (45.5 kt, 24 May 1951), an oralloy design exploded on island Janet at Enewetak. This experimental device used cryogenic liquid deuterium-tritium instead of gas. The boosting approximately doubled the yield over the expected unboosted value. Variants on the basic boosting approach that have been tested including the use of deuterium gas only, and the use of lithium deuteride/tritide, but it isn't known whether any of these approaches have been used in operational weapons.

Due to the marked increase in yield (as well as other reasons - such as reducing the weight of the fission system, and eliminating the risk of predetonation) today most fission bombs are boosted, including those used as primaries in true fission-fusion weapons. Although boosting can multiply the yield of fission bombs, it still has the same fundamental fission bomb design problems for high yield designs. The boosting technique is most valuable in small light-weight bombs that would otherwise have low efficiency. Tritium is a very expensive material to make, and it decays at a rate of 5.5% per year, but the small amounts required for boosting (a few grams) make its use economical.

1.5.2 Staged Radiation Implosion Weapons

This class of weapons is also called "Teller-Ulam" weapons, or (depending on type) fission-fusion or fission-fusion-fission weapons. These weapons use fusion reactions involving isotopes of light elements (e.g. hydrogen and lithium) to remove the yield limits of fission and boosted fission designs, to reduce weapon cost by reducing the amount of costly enriched uranium or plutonium required for a given yield, and to reduce the weight of the bomb. The fusion reactions occur in a package of fusion fuel ("the secondary") that is physically separate from the fission trigger ("the primary"), thus creating a two-stage bomb (the fission primary counting as the first stage). X-rays from the primary are used to compress the secondary through a process known as radiation implosion. The secondary is then ignited by a fission "spark plug" in its center. The energy produced by the fusion second stage can be used to ignite an even larger fusion third stage. Multiple staging allows in principle the creation of bombs of virtually unlimited size.

The fusion reactions are used to boost the yield in two ways:

1. By directly releasing a large amount of energy in fusion reactions;
2. By using high-energy or "fast" neutrons generated by fusion to release energy through fissioning of a fissionable jacket around the fusion stage. In the past this jacket was often made of natural or depleted uranium, so that energy is produced by fast fission of cheap U-238. Thorium may also be used for this purpose, although it is inferior to cheap and abundant depleted uranium. Weapons in which there is a premium on weight and size (i.e. virtually all modern strategic weapons) use moderately to highly enriched uranium as the jacketing material.

Bombs that release a significant amount of energy directly by fusion, but do not use fusion neutrons to fission the fusion stage jacket, are called Fission-Fusion weapons. If they employ the additional step of jacket fissioning using fusion neutrons they are called Fission-Fusion-Fission weapons.

The fast fission of the secondary jacket in a fission-fusion-fission bomb is sometimes thought of, or referred to, as a "third stage" in the bomb, and it is in a sense. But care must be taken not to confuse this with the true three-stage thermonuclear design in which there is another complete tertiary fusion stage.

Bombs that are billed as "clean" bombs (a relative term) obtain a large majority of their total yield from fusion. The last and largest stage of these bombs is always a pure fusion stage (not counting the spark plug), substituting a non-fissionable material for the jacket. The fusion-fraction of these designs as demonstrated in tests has been as high as 97% (this was the Tsar Bomba, see below).

Fission-fusion-fission bomb are dirty, but they have superior "bang for the buck" and "pow per pound". They generate a large amount of fission fallout since fission accounts for most of their yield. The 5 Mt Redwing Tewa test (20 July 1956 GMT, Bikini Atoll) had a fission fraction of 85%. If the emphasis is on cheapness depleted or natural uranium is usually used for the jacket. If the emphasis is on yield per weight (like nearly all modern strategic weapons) enriched uranium is used.

The staging concept allows the use as fuel pure deuterium, or varying mixtures of lithium 6 and 7 in the form of a compound with deuterium (lithium 6/7 deuteride). These natural stable isotopes are vastly cheaper than the artificially made and radioactive tritium.

The staged radiation implosion concept was originally conceived by Stanislaw Ulam and then developed further in a collaboration between Ulam and Edward Teller in early 1950. The first test of a staged thermonuclear device was Ivy Mike on 31 October 1952 (GMT) on Elugelab/Flora island at Enewetak Atoll. This experimental device, called Sausage, used pure deuterium fuel (probably the only time this was ever done) and a natural uranium jacket. It was designed by the Panda Committee led by J. Carson Mark at Los Alamos. Mike a yield of 10.4 Mt, 77% of which was fission.

The Teller-Ulam concept was later rediscovered by the other four nuclear weapon states, all of which have tested and deployed these weapons. No other nation is known to have deployed these designs, although the undeclared nuclear powers of Israel and India almost certainly have done development work on them.

Three stage designs have been tested and deployed to produce very high yield weapons. The first three stage U.S. test, and probably the first three stage weapon test ever, was the Bassoon device detonated in the Redwing Zuni test (27 May 1956 GMT, Bikini Atoll, 3.5 Mt). The largest nuclear explosion ever set off (50 Mt) was the Tsar Bomba (King of Bombs), a Soviet three

stage fission-fusion-fission design. It was exploded on 30 October 1961 over Novaya Zemlya at an altitude of 4000 m.

By jacketing the third stage with non-fissionable material, three stage devices can produce high yield clean weapons. Both Zuni and Tsar Bomba were in fact very clean devices - Zuni was 85% fusion and Tsar Bomba was 97% fusion. Both designs permitted replacing the lead or tungsten third stage jacket with U-238 however. A version of Bassoon called Bassoon Prime was tested in the dirty Tewa test mentioned above. A dirty device derived from the Bassoon was weaponized to create the highest yield weapon the U.S. ever fielded, the 25 megaton Mk-41. The Tsar Bomba design was for a fission-fusion-fission bomb with a staggering yield of at least **100 megatons**!

A possible variation on the staged radiation implosion design is one in which a second fission stage is imploded instead of a thermonuclear one. This was actually the initial concept developed by Stanislaw Ulam before he realized its possible application to thermnuclear weapons. The advantage of this approach is that radiation implosion speeds are hundreds of times higher, and maximum densities tens of times greater, than those achievable through high explosives. This allows achieving higher yields than is practical with high explosive driven fission weapons, and the use of lower grades of fissile material. If some fusion fuel is included in this second fission stage to boost yield, a sort of hybrid two-stage boosted weapon design results that blurs the distinction between two-stage fission and classic Teller-Ulam thermonuclear weapons. The TX-15 "Zombie" developed by the U.S. was originally planned to be a two stage pure fission device, but later evolved into this sort of hybrid boosted system. The Zombie was tested in the Castle Nectar shot (13 May 1954 GMT; Bikini Atoll; 1.69 Mt), and was fielded as the Mk-15.

1.5.3 The Alarm Clock/Sloika (Layer Cake) Design

This idea predates the invention of staged radiation implosion designs, and was apparently invented independently at least three times. It was first devised by Edward Teller in the United States, who named the design "Alarm Clock". Later Andrei Sakharov and Vitalii Ginzburg in the Soviet Union hit upon it and dubbed it the "sloika" design. A sloika is a layered Russian pastry, rather like a napoleon, and has thus been translated as "Layer Cake". Finally it was developed by the British (inventor unknown). Each of these weapons research programs hit upon this idea before ultimately arriving at the more difficult, but more powerful and efficient, staged thermonuclear approach.

This system was dubbed "Layer Cake" by the Soviets because it uses a spherical assembly of concentric shells. In the center is a fission primary made of U-235/Pu-239, surrounding it is an (optional) layer of U-238 for the fission tamper, then a layer of lithium-6 deuteride/tritide, a U-238 fusion tamper, and finally the high explosive implosion system. The process begins like an ordinary implosion bomb. After the primary in the center completes its reaction, the energy it releases compresses and heats the fusion layer to thermonuclear temperatures. The burst of fission neutrons then initiates a coupled fission-fusion-fission chain reaction. Slower fission neutrons generate tritium from the lithium, which then fuses with deuterium to produce very fast

neutrons, that in turn cause fast fission in the fusion tamper, which breed more tritium. In effect the fusion fuel acted as a neutron accelerator allowing a fission chain reaction to occur with a large normally non-fissionable U-238 mass. While spiking the fusion layer with an initial dose of tritium is not strictly necessary for this approach, it helps boost the yield.

The achievable fusion fraction is fairly small, 15-20%, and cannot be increased beyond this point. Its use of fusion fuel is also quite inefficient. This design is also limited to the same yield and yield-to-weight range as high yield pure fission and boosted fission weapons. This was developed into a deliverable weapon by the Soviet Union and the British prior to their development of the staged designs described above. The U.S. did not bother to pursue it, partly because Teller did not feel it was sufficiently destructive to be worthwhile.

The first test of this concept was a device designated RDS-6s, (known as Joe 4 to the U.S.) on 12 August 1953. By using tritium doping it achieved a 10-fold boost over the size of the trigger, for a total yield of 400 kt. The UK Orange Herald Small device tested in Grapple 2 (31 May 1957) was similar but used a much larger fission trigger (300 kt range) apparently without tritium for a total yield of 720 kt, a boost in the order of 2.5-fold. This is probably the largest test of this design.

Although apparently not used in any weapons now in service in the five declared weapons states, it remains a viable design that could be attractive to other states that do not have the resources to develop the technically more demanding radiation implosion design. Information supplied by Mordechai Vanunu indicates that Israel may have developed a weapon of this type.

This design should probably be considered distinct from other classes of nuclear weapons. This design is something of a hybrid and could be considered either a type of boosted fission device, or a one-stage type of fission-fusion-fission bomb.

1.5.4 Neutron Bombs

Neutron bombs, more formally referred to as "enhanced radiation (ER) warheads", are small thermonuclear weapons in which the burst of neutrons generated by the fusion reaction is intentionally not absorbed inside the weapon, but allowed to escape. This intense burst of high-energy neutrons is the principle destructive mechanism. Neutrons are more penetrating than other types of radiation so many shielding materials that work well against gamma rays do not work nearly as well. The term "enhanced radiation" refers only to the burst of ionizing radiation released at the moment of detonation, not to any enhancement of residual radiation in fallout.

The U.S. has developed neutron bombs for use as strategic anti-missile weapons, and as tactical weapons intended for use against armored forces. As an anti-missile weapon ER weapons were developed to protect U.S. ICBM silos from incoming Soviet warheads by damaging the nuclear components of the incoming warhead with the intense neutron flux. Tactical neutron bombs are primarily intended to kill soldiers who are protected by armor. Armored vehicles are extremely resistant to blast and heat produced by nuclear weapons, so the effective range of a nuclear weapon against tanks is determined by the lethal range of the

radiation, although this is also reduced by the armor. By emitting large amounts of lethal radiation of the most penetrating kind, ER warheads maximize the lethal range of a given yield of nuclear warhead against armored targets.

One problem with using radiation as a tactical anti-personnel weapon is that to bring about rapid incapacitation of the target, a radiation dose that is many times the lethal level must be administered. A radiation dose of 600 rads is normally considered lethal (it will kill at least half of those who are exposed to it), but no effect is noticeable for several hours. Neutron bombs were intended to deliver a dose of 8000 rads to produce immediate and permanent incapacitation. A 1 kt ER warhead can do this to a T-72 tank crew at a range of 690 m, compared to 360 m for a pure fission bomb. For a "mere" 600 rad dose the distances are 1100 m and 700 m respectively, and for unprotected soldiers 600 rad exposures occur at 1350 m and 900 m. The lethal range for tactical neutron bombs exceeds the lethal range for blast and heat even for unprotected troops.

The neutron flux can induce significant amounts of short lived secondary radioactivity in the environment in the high flux region near the burst point. The alloy steels used in armor can develop radioactivity that is dangerous for 24-48 hours. If a tank exposed to a 1 kt neutron bomb at 690 m (the effective range for immediate crew incapacitation) is immediately occupied by a new crew, they will receive a lethal dose of radiation within 24 hours.

Newer armor designs afford more protection than the Soviet T-72 against with ER warheads were initially targeted. Special neutron absorbing armor techniques have also been developed and deployed, such as armors containing boronated plastics and the use of vehicle fuel as a shield. Some newer types of armor, like that of the M-1 tank, employ depleted uranium which can offset these improvements since it undergoes fast fission, generating additional neutrons and becoming radioactive.

Due to the rapid attenuation of neutron energy by the atmosphere (it drops by a factor of 10 every 500 m in addition to the effects of spreading) ER weapons are only effective at short ranges, and thus are found in relatively low yields. ER warheads are also designed to minimize the amount of fission energy and blast effect produced relative to the neutron yield. The principal reason for this was to allow their use close to friendly forces. The common perception of the neutron bomb as a "landlord bomb" that would kill people but leave buildings undamaged is greatly overstated. At the intended effective combat range (690 m) the blast from a 1 kt neutron bomb will destroy or damage to the point of unusability almost any civilian building. Thus the use of neutron bombs to stop an enemy attack, which requires exploding large numbers of them to blanket the enemy forces, would also destroy all buildings in the area.

Neutron bombs (the tactical versions at least) differ from other thermonuclear weapons in that a deuterium-tritium gas mixture is the only fusion fuel. The reasons are two-fold: the D-T thermonuclear reaction releases 80% of its energy as neutron kinetic energy, and it is also the easiest of all fusion reactions to ignite. This means that only 20% of the fusion energy is available for blast and thermal radiaiton production, that the neutron flux produced consists of extremely penetrating 14.7 Mev neutrons, and that a very small fission explosion (250-400 tons)

can be used for igniting the reaction. The more typical lithium deuteride fuel would produce much more blast and flash for each unit of neutron flux, and would require a much larger fission explosion to set it off. The disadvantage of using D-T fuel is that tritium is very expensive, and decays at a rate of 5.5% a year. Combined with its increased complexity this makes ER warheads more expensive to build and maintain than other tactical nuclear weapons. To produce a 1 kt fusion yield 12.5 g of tritium and 5 g of deuterium are required.

The U.S. developed and produced three neutron warheads, a fourth was cancelled prior to production. All have been retired and dismantled.

- The W66 warhead for the Sprint missile was the first ER warhead to be developed. It was manufactured during 1974-75, and was retired in Aug. 1975 after only a few months of service when the Sprint system was deactivated (about 70 were made). It had a yield of several kilotons (20 kt has been reported) and may or may not have used D-T fuel.
- The W70 Mod 3 warhead for the Lance missile had a total yield about 1 kt which was 60% fusion and 40% fission. It was manufactured during 1981-83, and was retired by 1992; 380 were built.
- The W79 Mod 0 warhead for the 8 inch artillery shell had a variable yield from 100 T to 1.1 kt. At the lowest yield it was a pure fission weapon, at the highest yield 800 T was from fusion (73%) and 300 T from fission. It was manufactured during 1981-1986; this version began retirement in the mid-80s, all were retired by 1992; 325 were built.
- The W82 Mod 0 155 mm artillery shell, with variable yields similar tot he W79, was canceled in Oct 1983 without going in to production.

The Soviet Union, China, and France are all known to have developed neutron bomb designs and may have them in service. A number of reports have claimed that Israel has developed neutron bombs which, though they could be valuable on an armor battleground like the Golan Heights, are difficult to develop and require sigificant testing. This makes it unlikely that Israel has in fact acquired them.

1.6 Cobalt Bombs and other Salted Bombs

A "salted" nuclear weapon is reminiscent of fission-fusion-fission weapons, but instead of a fissionable jacket around the secondary stage fusion fuel, a non-fissionable blanket of a specially chosen salting isotope is used (cobalt-59 in the case of the cobalt bomb). This blanket captures the escaping fusion neutrons to breed a radioactive isotope that maximizes the fallout hazard from the weapon rather than generating additional explosive force (and dangerous fission fallout) from fast fission of U-238.

Variable fallout effects can be obtained by using different salting isotopes. Gold has been proposed for short-term fallout (days), tantalum and zinc for fallout of intermediate duration (months), and cobalt for long term contamination (years). To be useful for salting, the parent

isotopes must be abundant in the natural element, and the neutron-bred radioactive product must be a strong emitter of penetrating gamma rays.

Table 1.6-1 Candidate Salting Agents

Parent Isotope	Natural Abundance	Radioactive Product	Half-Life
Cobalt-59	100%	Co-60	5.26 years
Gold-197	100%	Au-198	2.697 days
Tantalum-181	99.99%	Ta-182	115 days
Zinc-64	48.89%	Zn-65	244 days

The idea of the cobalt bomb originated with Leo Szilard who publicized it in Feb. 1950, not as a serious proposal for weapon, but to point out that it would soon be possible in principle to build a weapon that could kill everybody on earth (see Doomsday Device in Questions and Answers). To design such a theoretical weapon a radioactive isotope is needed that can be dispersed world wide before it decays. Such dispersal takes many months to a few years so the half-life of Co-60 is ideal.

The Co-60 fallout hazard is greater than the fission products from a U-238 blanket because

1. many fission-produced isotopes have half-lives that are very short, and thus decay before the fallout settles or can be protected against by short-term sheltering;
2. many fission-produced isotopes have very long half-lives and thus do not produce very intense radiation;
3. the fission products are not radioactive at all. The half-life of Co-60 on the other hand is long enough to settle out before significant decay has occurred, and to make it impractical to wait out in shelters, yet is short enough that intense radiation is produced. Initially gamma radiation fission products from an equivalent size fission-fusion-fission bomb are much more intense than Co-60: 15,000 times more intense at 1 hour; 35 times more intense at 1 week; 5 times more intense at 1 month; and about equal at 6 months. Thereafter fission drops off rapidly so that Co-60 fallout is 8 times more intense than fission at 1 year and 150 times more intense at 5 years. The very long lived isotopes produced by fission would overtake the again Co-60 after about 75 years.
   Zinc has been proposed as an alternate candidate for the "doomsday role". The advantage of Zn-64 is that its faster decay leads to greater initial intensity. Disadvantages are that since it makes up only half of natural zinc, it must either be isotopically enriched or the yield will be cut in half; that it is a weaker gamma emitter than Co-60, putting out only one-fourth as many gammas for the same molar quantity; and that substantially amounts will decay during the world-wide dispersal process. Assuming pure Zn-64 is used, the radiation intensity of Zn-65 would initially be twice as much as Co-60. This would decline to being equal in 8 months, in 5 years Co-60 would be 110 times as intense.

Militarily useful radiological weapons would use local (as opposed to world-wide) contamination, and high initial intensities for rapid effects. Prolonged contamination is also undesirable. In this light Zn-64 is possibly better suited to military applications than cobalt, but probably inferior to tantalum or gold. As noted above ordinary "dirty" fusion-fission bombs have very high initial radiation intensities and must also be considered radiological weapons.

No cobalt or other salted bomb has ever been atmospherically tested, and as far as is publicly known none have ever been built. In light of the ready availability of fission-fusion-fission bombs, it is unlikely any special-purpose fallout contamination weapon will ever be developed.

The British did test a bomb that incorporated cobalt as an experimental radiochemical tracer (Antler/Round 1, 14 September 1957). This 1 kt device was exploded at the Tadje site, Maralinga range, Australia. The experiment was regarded as a failure and not repeated.

2.0 Introduction to Nuclear Weapon Physics and Design

Discussions of physical principle, particularly nuclear physics, is unavoidable in most of the sections of this FAQ. In this section I set forth the basic principles behind all nuclear weapons, although some familiarity with physics is assumed. Section 4 deals with the design and engineering of nuclear weapons in more detail, and the physics discussions there can be considered a continuation of Section 2.

- 2.1 Fission Weapon Physics
- 2.1.1 The Nature Of The Fission Process
- 2.1.2 Criticality
- 2.1.3 Time Scale of the Fission Reaction
- 2.1.4 Basic Principles of Fission Weapon Design
- 2.2 Fusion Weapon Physics
- 2.2.1 Candidate Fusion Reactions
- 2.2.2 Basic Principles of Fusion Weapon Design

2.1 Fission Weapon Physics

Nuclear fission occurs when the nuclei of certain isotopes of very heavy elements, isotopes of uranium and plutonium for example, capture neutrons. The nuclei of these isotopes are just barely stable and the addition of a small amount of energy to one by an outside neutron will cause it to promptly split into two roughly equal pieces, with the release of a great deal of energy (180 MeV of immediately available energy) and several new neutrons (an average of 2.52 for U-235, and 2.95 for Pu-239). If on average one neutron from each fission is captured and successfully produces fission then a self-sustaining chain reaction is produced. If on

average *more* than one neutron from each fission triggers another fission, then the number of neutrons and the rate of energy production will increase exponentially with time.

Two conditions must be met before fission can be used to create powerful explosions: 1) the number of neutrons lost to fission (from non-fission producing neutron captures, or escape from the fissionable mass) must be kept low, and 2) the speed with which the chain reaction proceeds must be very fast. A fission bomb is in a race with itself: to successfully fission most of the material in the bomb before it blows itself apart. The degree to which a bomb design succeeds in this race determines its efficiency. A poorly designed or malfunctioning bomb may "fizzle" and release only a tiny fraction of its potential energy.

2.1.1 The Nature Of The Fission Process

The nucleus of an atom can interact with a neutron that travels nearby in two basic ways. It can scatter the neutron - deflecting the neutron in a different direction while robbing it of some of its kinetic energy. Or it can capture the neutron, which in turn can affect the nucleus in several ways - absorption and fission being most important here. The probability that a particular nucleus will scatter or capture a neutron is measured by its scattering cross-section and capture cross-section respectively. The overall capture cross-section can be subdivided into other cross-sections - the absorption cross-section and the fission cross-section.

The stability of an atomic nucleus is determined by its binding energy - the amount of energy required to disrupt it. Any time a neutron or proton is captured by an atomic nucleus, the nucleus rearranges its structure. If energy is released by the rearrangement, the binding energy decreases. If energy is absorbed, the binding energy increases.

The isotopes important for the large scale release of energy through fission are uranium-235 (U-235), plutonium-239 (Pu-239), and uranium-233 (U-233). The binding energy of these three isotopes is so low that when a neutron is captured, the energy released by rearrangement exceeds it. The nucleus is then no longer stable and must either shed the excess energy, or split into two pieces. Since fission occurs regardless of the neutron's kinetic energy (i.e. no extra energy from its motion is needed to disrupt the nucleus), this is called "slow fission".

By contrast, when the abundant isotope uranium-238 captures a neutron it still has a binding energy deficit of 1 MeV after internal rearrangement. If it captures a neutron with a kinetic energy exceeding 1 MeV, then this energy plus the energy released by rearrangement can over come the binding energy and cause fission. Since a fast neutron with a large kinetic energy is required, this is called "fast fission".

The slow fissionable isotopes have high neutron fission cross-sections for neutrons of all energies, while having low cross-sections for absorption. Fast fissionable isotopes have zero fission cross-sections below a certain threshold (1 MeV for U-238), but the cross-sections climb

quickly above the threshold. Generally though, slow-fissionable isotopes are more fissionable than fast-fissionable isotopes for neutrons of all energies.

A general trend among the elements is that the ratio of neutrons to protons in an atomic nucleus increases with the element's atomic number (the number of protons the nucleus contains, which determines which element it is). Heavier elements require relatively more neutrons to stabilize the nucleus. When the nucleus of a heavy element like uranium (atomic number 92) is split the fragments, having lower atomic numbers, will tend to have excess neutrons. These neutrons are shed very rapidly by the excited fragments. More neutrons are produced on average than are consumed in fission.

Fission is a statistical process. The nucleus rarely splits into pieces with nearly the same mass and atomic number. Instead both the size and atomic numbers of the fragments have a Gaussian distributions around two means (one for the lighter fragment around 95, one for the heavier around 135). Similarly, the number of neutrons produced varies from zero to six or more, and their kinetic energy varies from 0.5 MeV to more than 4 MeV, the most probable energy is 0.75 MeV, the average (and median) is 2 MeV.

A breakdown of the energy released by fission is given below:

	MeV
Kinetic energy of fission fragments	165 +/- 5
Instantaneous gamma rays	7 +/- 1
Kinetic energy of neutrons	5 +/- 0.5
Beta particles from product decay	7 +/- 1
Gamma rays from product decay	6 +/- 1
Neutrinos from product decay	10
TOTAL	200 +/- 6

All of the kinetic energy is released to the environment instantly, as are most of the instantaneous gamma rays. The neutrino energy effectively disappears; the unstable fission products release their decay energies at varying rates, some almost immediately. The net result is that about 180 MeV is actually available to generate nuclear explosions, the remainder of the decay energy shows up over time as fallout (or is carried away by the virtually undetectable neutrinos).

2.1.2 Criticality

A neutron entering a pure chunk of one of the slow-fissionable isotopes would have a high probability of causing fission compared with the chance of unproductive absorption. If the chunk is large and compact enough, then the rate of neutron escape from its surface will be so low that it becomes a "critical mass", a mass in which a self-sustaining chain reaction occurs. Non-

fissionable materials mixed with these isotopes tend to absorb some of the neutrons uselessly and also dilute it (decrease its effective density), both effects increase the required size of the critical mass or may even make it impossible to achieve altogether.

Typical figures for critical masses for bare (unreflected) spheres of fissionable materials are:

U-233	16 kg
U-235	48 kg
Pu-239 (alpha phase)	10 kg

2.1.3 Time Scale of the Fission Reaction

The amount of time taken by each link in the chain reaction is determined by the speed of the neutrons and the distance they travel before being captured. The average distance is called the mean free path. In fissile materials at maximum normal densities the mean free path for fission is roughly 13 cm for 1 MeV neutrons (a typical energy for fission neutrons). These neutrons travel at 1.4x10^9 cm/sec, yielding an average time between fission generations of about 10^-8 sec (10 nanoseconds), a unit of time sometimes called a "shake". The mean free path for scattering is only 2.5 cm, so on average a neutron will be scattered 5 times before causing fission.

Actual 1 MeV mean free path values are:

	Density	M.F.P. (cm)
U-233	18.9	10.9
U-235	18.9	16.5
Pu-239	19.4	12.7

This shows that fission proceeds faster in some isotopes than others.

The rate of multiplication can be calculated from the multiplication coefficient k given by:

$$k = f - (lc + le)$$

where f = avg. neutrons generated per fission
lc = avg. neutrons lost to capture
le = avg. neutrons lost by escaping assembly

When k = 1 an assembly is exactly critical and a chain reaction will be self supporting, although it will not increase in rate. When k 1 then it is super-critical and the reaction will continually increase. To make an efficient bomb k must be as high as possible, usually somewhere near 2, when the chain reaction starts.

Many discussions of fission describe the chain reaction as proceeding by discrete generations. Generation zero has 1 neutron, generation one has 2 neutrons, generation two has 4 neutrons, etc. until, say, 2x10^24 atoms have been split - which produces 20 kilotons of energy. The formula for this is:

Number of atoms split = 2^(n-1), where n is the generation number.

Thus 2x10^24 = 2^(n-1) implies n = (log2 (2x10^24)) + 1 = 81.7 generations. That is, it takes about 82 generations to complete the fission process for a 20 kiloton bomb, if the reaction starts from one neutron.

This calculation is a useful simplification, but the fission process does not really proceed by separate steps, each completing before the next begins. It is really a continuous process, the current oldest generation of neutrons starts creating the next generation even while it is still being formed by neutrons from still older generations. An accurate calculation thus requires the use of formulas derived from calculus.

We find that both the number of neutrons present in the assembly (and thus the instantaneous rate of the fission reaction), and the number of fissions that have occurred since the reaction began, increase at a rate proportional to e^((k-1)*(t/g)), where e is the natural log base (2.712...), g is the average generation time (time from neutron emission to fission capture), and t is the elapsed time.

If k=2, then a single neutron will multiply to 2x10^24 neutrons (and splitting the same number of atoms) in roughly 56 shakes (560 nanoseconds), yielding 20 kilotons of energy. This is one-third less time than the previous approximate calculation. Due to the exponential rate of increase, at any point in the chain reaction 99% of the energy will have been released in the last 4.6 generations. It is a reasonable approximation to think of the first 53 generations as a latency period leading up to the actual explosion, which only takes 3-4 generations.

The extremely rapid buildup in the fission rate as the reaction proceeds has some important consequences that should be pointed out. The longer a neutron takes to cause fission, the less significant it is in contributing to the chain reaction. This is because it becomes quickly outnumbered by the descendants of neutrons that undergo fission capture sooner. Thus faster, more energetic, neutrons contribute disproportionately compared to slower neutrons. This is called "time absorption" since it has the same effect as a neutron absorber with a cross-section inversely proportional to velocity. Similarly, if a neutron leaves the critical mass and is scattered back in, then its contribution is also considerably reduced. In fact since the path of a neutron that leaves the critical assembly, then re-enters is much longer than the average path of neutron that remains within the mass the time absorption in a reflector is very large.

2.1.4 Basic Principles of Fission Weapon Design

The principle issues that must solved to construct a fission weapon are:

1. Keeping the fissionable material in a subcritical state before detonation;
2. Bringing the fissionable material into a supercritical mass while keeping it free of neutrons;
3. Introducing neutrons into the critical mass when it is at the optimum configuration (i.e. at maximum supercriticality);

Keeping the mass together until a substantial portion of the material has fissioned.

Solving issues 1, 2 and 3 together is greatly complicated by the unavoidable presence of naturally occurring neutrons. Although cosmic rays generate neutrons at a low rate, almost all of these "background" neutrons originate from the fissionable material itself through the process of spontaneous fission. Due to the low stability of the nuclei of fissionable elements, these nuclei will occasionally split without being hit by a neutron. This means that the fissionable material itself periodically emits neutrons.

The process of assembling the supercritical mass must occur in significantly less time than the average interval between spontaneous fissions to have a reasonable chance of succeeding. This problem is difficult to accomplish due to the very large change in reactivity required in going from a subcritical state to a supercritical one. The time required to raise the value of k from 1 to the maximum value of 2 or so is called the reactivity insertion time, or simply insertion time.

It is further complicated by the problem of subcritical neutron multiplication. If a subcritical mass has a k value of 0.9, then a neutron present in the mass will (on average) create a chain reaction that dies out in an average of 10 generations. If the mass is very close to critical, say k=0.99, then each spontaneous fission neutron will create a chain that lasts 100 generations. This persistence of neutrons in subcritical masses further reduces the time window for assembly, and requires that the reactivity of the mass be increased from a value of less than 0.9 to a value of 2 or so within that window.

Simply splitting a supercritical mass into two identical parts, and bringing the parts together rapidly is unlikely to succeed since neither part will have a sufficiently low k value, nor will the insertion time be rapid enough with achievable assembly speeds.

2.1.4.1 Assembly Techniques - Achieving Supercriticality

The key to achieving objectives 1 and 2 is revealed by the fact that the critical mass (or supercritical mass) of a fissionable material is inversely proportional to the square of its density. By contriving a subcritical arrangement of fissionable material whose average density can be rapidly increased, we can bring about the sudden large increase in reactivity needed to create a powerful explosion. As a general guide, a suitable highly supercritical mass needs to be at least three times heavier than a mass of equal density and shape that is merely critical. Thus doubling the density of a pit that is slightly sub-critical (thereby making it into nearly four critical masses) provides sufficient reactivity insertion for a bomb.

Two general approaches have been used for achieving this idea: implosion assembly, and gun assembly. Implosion is capable of very short insertion times, gun assembly is much slower.

2.1.4.1.1 Implosion Assembly

The key idea in implosion assembly is to compress a subcritical spherical, or sometimes cylindrical, fissionable mass by using specially designed high explosives. Implosion works by initiating the detonation of the explosives on their outer surface, so that the detonation wave moves inward. Careful design allows the creation of a smooth, symmetrical implosion shock wave. This shock wave is transmitted to the fissionable core and compresses it, raising the density to the point of supercriticality.

Implosion can be used to compress either solid cores of fissionable material, or hollow cores in which the fissionable material forms a shell. It is easy to see how implosion can increase the density of a hollow core - it simply collapses the cavity. Solid metals can be compressed substantially by powerful shock waves also though. A high performance explosive can generate shock wave pressures of 400 kilobars (four hundred thousand atmospheres), implosion convergence and other concentration techniques can boost this to several megabars. This pressure can squeeze atom closer together and boost density to twice normal or even more (the theoretical limit for a shock wave in an ideal monatomic gas is a four-fold compression, the practical limit is always lower).

The convergent shock wave of an implosion can compress solid uranium or plutonium by a factor of 2 to 3. The compression occurs very rapidly, typically providing insertion times in the range to 1 to 4 microseconds. The period of maximum compression lasts less than a microsecond.

A two-fold compression will boost a slightly sub-critical solid mass to nearly four critical masses. Such a solid core design was used for Gadget, the first nuclear explosive ever tested, and Fat Man, the atomic bomb dropped on Nagasaki. In practice hollow core designs also achieve greater than normal densities (i.e. they don't rely on collapsing a hollow core alone).

In addition to its major objective of achieving supercriticality, compression has another important effect. The increased density reduces the neutron mean free path, which is inversely proportional to density. This reduces the time period for each generation and allows a faster reaction that can progress farther before disassembly occurs. Implosion thus considerably increases a bomb's efficiency.

The primary advantages of implosion are:

a. high insertion speed - this allows materials with high spontaneous fission rates (i.e. plutonium) to be used;
b. high density achieved, leading to a very efficient bomb, and allows bombs to be made with relatively small amounts of material;
c. potential for light weight designs - in the best designs only several kilograms of explosive are needed to compress the core.

The principal drawback is its complexity and the precision required to make it work. Implosion designs take extensive research and testing, and require high precision machining and electronics.

2.1.4.1.2 Gun Assembly

Assembling a critical mass by firing one piece of fissionable material at another is an obvious idea and was the first approach developed for designing atomic bombs. But it is probably not obvious how you take two subcritical masses and obtain the equivalent of three critical masses by bringing them together.

This can be made clear by conducting a thought-experiment. Imagine a spherical pit made up of about three critical masses of fissionable material. Now remove a core (like an apple core) from the pit with a mass slightly less than critical. Since the center of the pit is now hollow, its effective density has been reduced to 2/3 of the original density. Since we now have two critical masses remaining in the pit, and the reduction in density leads to a further reduction of (2/3)^2 = 4/9, the pit now contains only 2*(4/9) = 8/9 of a critical mass.

The two sub-critical pieces can be brought together by firing the cylindrical core down a gun barrel into the center of the hollowed-out pit. The insertion time is large - over 1 millisecond. This is the design used in Little Boy, the bomb dropped on Hiroshima (except that a slightly less efficient squat cylinder was used, rather than a spherical pit).

The primary advantage of gun assembly is simplicity. It is as close to a fool-proof design as ordinance technology allows.

The drawbacks are:

a. the lack of compression, which requires large amounts of fissionable material, and leads to low efficiency;
b. only uranium-235 (and possibly U-233) can be used due to the slow insertion speed;
c. the weight and length of the gun barrel makes the weapon heavy and fairly long.

2.1.4.2 Initiating Fission

Assembly techniques only address issues 1 and 2, reconfiguring sub-critical masses rapidly into supercritical ones. The next problem is to make sure fission does occur when it is desired.

Since neutrons are generated periodically by spontaneous fission, one approach would be to hold the supercritical mass together after it is assembled until spontaneous neutrons start the reaction. This is at least possible for gun assembly, but it is unsatisfactory for implosion since the highly compressed pit begins expanding soon after the shock wave dies out. Even in a compressed pit the fission reaction takes about 250 nanoseconds, roughly the duration of the maximum compression. It is therefore important to initiate the chain reaction very soon after maximum compression is achieved, or even slightly before.

A better method is to have some sort of neutron generator whose operation is precisely synchronized with the assembly process. Three general mechanisms have been developed for this, all of which use charged particle reactions to generate neutrons.

The first type of generator to be invented relies on the fact that one of the neutrons in beryllium-9 is easily knocked loose. Occasionally if it is struck by an alpha particle, like those produced by some produced by some radioactive isotopes, a neutron will be released as a result of the collision:

Be-9 + He-4 -> Be-8 + n + He-4

This happens in only 0.008% of collisions, so a strong alpha emitter (like polonium-210) is required to achieve the neutron flux needed by an implosion weapon. A neutron generation rate of 10-100 million neutrons per second is needed to ensure the prompt initiation of the reaction, thus 100-1000 billion alphas per second are required (3-30 curies of radioactive material). The generator is located in the center of the pit. Clever designs (still classified in the US, though detailed descriptions now exist in the open literature) are needed to keep the alpha emitter and beryllium separate, but still allow the implosion process to bring them together rapidly. This type of generator was used in all of the early atomic weapons.

The major problem with the beryllium/alpha emitter generators is that the strong emitters used have very short half-lifes (138.4 days for Po-210). Maintaining a inventory of weapons thus requires continual manufacture and replacement of generators. Also, due to difficulties in precisely controlling the mixing of the beryllium and polonium it is difficult to control the initiation of the fission reaction accurately. These types of generators had a tendency to start the reaction later than optimum.

A somewhat similar approach is to use the implosion to initiate a neutron generating fusion reactions with tritium and deuterium (described in Section 2.2 below). It may seem surprising that this can be made to work, given the well known fact that fission explosions are required to produce the temperatures that fusion reactions normally need. Three considerations overcome this obstacle. First, an exceedingly low rate of fusion is actually required. One neutron (and thus one fusion) every 10 nanoseconds is sufficient, a rate that is only some 10^-24 as fast as an actual fusion explosion would need. Second, implosions focus energy and can reach very high temperatures near the center. Theoretically the temperature at the center is infinitely high, but lack of perfect symmetry reduces this. Even so, a high precision implosion can reach temperatures of several hundred thousand degrees C. Third, the velocity of atoms in a gas or plasma is a statistical (Maxwellian) distribution. A very small portion of the atoms can greatly exceed the average energy. Thus enough atoms in the D-T mixture near the center can reach fusion energies to produce the required rate of neutron production. This type of implosion initiator is even more difficult to engineer than the Be/Po-210 type since the very high precision implosion is required to achieve the required symmetry. The major advantage is that the short half-life Po-210 is not needed.

A more sophisticated system is to use an electronically controlled particle accelerator called a pulse neutron tube. These generators use the deuterium+deuterium or deuterium+tritium fusion reactions to produce large amounts of neutrons. A very short surge of high voltage current accelerates a pulse deuterium or tritium nuclei to energies sufficient to cause fusion reactions, then slams them into a deuterium or tritium rich target. A short pulse containing millions of neutrons is produced. The timing of the pulse can be precisely controlled. Because of the large number of neutrons produced, the generator can be located anywhere in the weapon with assurance that a sufficient number will reach the pit. This is the initiator commonly used in most modern nuclear weapons.

2.1.4.3 Preventing Disassembly and Increasing Efficiency

By the time a significant percentage of the atoms have fissioned, their thermal kinetic energy is so high that the pit will expand enough to shut down the reaction in only a few of shakes. This severely limits the efficiency of fission weapons (percentage of material fissioned). The practical efficiency limit of a typical pure fission bomb is about 25%, and could be much less. The Fat Man implosion bomb was 17% efficient (counting only the energy produced by the fissile core, the natural uranium tamper contributed another 4% through fast fission). Little Boy had an efficiency of only 1.4%. Very large pure fission bombs can achieve efficiencies approaching 50% but have been supplanted by thermonuclear weapon technology. Anything that will increase the confinement time of the fissionable core or decrease the generation time, even slightly, can cause a significant increase in bomb yield.

As noted above, compressing the fissionable material through implosion makes the largest contribution to a bomb's efficiency. By doubling the density of the pit the length of a shake is cut in half, thus allowing almost twice as many fission generations to occur during the brief period before expansion halts the reaction.

Another approach for increasing efficiency is reducing the rate of expansion through better confinement of the critical mass. A layer of dense material called a "tamper" (typically made of natural or depleted uranium or tungsten) surrounds the critical mass. A bare hot critical mass does not expand uniformly. The material in the center of the mass is confined by the pressure of the outer layers of the mass, and thus does not expand initially. In the absence of a tamper the outer surface of the mass has no external pressure holding it in, and thus begins rapid expansion immediately. The material blows off at supersonic speeds, and an expansion wave moves inward at the speed of sound.

A tamper improves confinement for two reasons. First, the expanding material must drive a shock wave through the dense tamper rather than expanding into a vacuum. This dramatically reduces the rate of expansion. Second, a layer of the tamper next to the mass is heated by the explosion and exerts pressure on the surface holding it together. This delays disassembly further since the expansion wave must first traverse this hot tamper layer before the fissile material can begin to expand.

The tamper has an additional benefit, it can also scatter or "reflect" neutrons back into the critical mass after they escape from its surface. This means that a smaller amount of fissionable material is necessary to make the critical mass. The importance of this effect is often overstated in the nuclear weapons literature however. Only a portion of the neutrons are scattered back, and since it takes on average several shakes for the neutrons that do return to reenter the critical mass, their significance is further reduced through "time absorption" (see section 2.1.3). This is offset somewhat by the fact that some neutron multiplication occurs in natural uranium tampers through fast fission of U-238.

2.2 Fusion Weapon Physics

Fusion reactions, also called thermonuclear reactions, are reactions between the nuclei of certain isotopes of light elements. If the nuclei collide with sufficient energy (provided by heat in a bomb, or by a particle accelerator in the laboratory) then there is a significant chance that they will merge to form one or more new nuclei with the release of energy. Different nuclei combinations have different inherent likelihoods of reacting in a collision at a particular temperature. The rates of all fusion reactions is affected by both temperature and density. The hotter and denser the fusion fuel, the faster the fusion "burn".

The fusion reactions that occur in stars are not the same as the ones that occur in thermonuclear weapons or (laboratory fusion reactors). The somewhat complex catalyzed fusion cycle in stars that converts light hydrogen (protium) into helium is extremely slow, which is why the lifetime of the Sun is measured in billions of years. The fusion reactions used in bombs and prospective powerplant designs are simple, and extremely fast - which is essential since the fuel must be fully consumed within microseconds. These reactions thus are based on the same general principles as stellar fusion, but are completely different in detail.

2.2.1 Candidate Fusion Reactions

The most important fusion reactions for thermonuclear weapons are given below:

1. D + T - He-4 + n + 17.588 MeV
2. D + D - He-3 + n + 3.268 MeV
3. D + D - T + p + 4.03 MeV
4. He-3 + D - He-4 + p + 18.34 MeV
5. Li-6 + n - T + He-4 + 4.78 MeV
6. Li-7 + n - T + He-4 + n - 2.47 MeV

[D and T stand for deuteron or deuterium (H-2), and triton or tritium (H-3) respectively.]

At the temperatures found in fission bombs reaction 1 has a rate 100 times faster than the next fastest candidate (reactions 2 and 3 combined), which are in turn 10 times faster than reaction 4. The rates of reactions 1 - 4 all increase rapidly (exponentially) with temperature, but not in the same proportion. At the higher temperatures achievable by fusion, reaction 4 exceeds the combined rate of reactions 2 and 3. Other reactions also occur between the isotopes listed here, but the reactions rates are too low to be important.

Some additional important facts about these reactions:

The neutron produced in reaction 1 is extremely energetic, it carries away 14.06 MeV of the reaction energy, the alpha particle (He-4 nucleus) only 3.52 MeV.

The neutron produced in reaction 2 has an energy of only 2.45 MeV (similar to the faster fission neutrons), with the He-3 carrying 0.82 MeV. The division of energy in reaction 3 is 1.01 MeV for the triton, and 3.03 MeV for the proton. The two D+D reactions are equally likely and each will occur half the time.

In reaction 4 the alpha particle carries off 3.67 MeV, the proton 14.67 MeV.

Reactions 5 and 6 are not thermonuclear reactions, strictly speaking. They are neutronic reactions, like fission, and do not require heat or pressure, just neutrons in the correct energy range. This distinction is usually ignored in the literature about nuclear weapons however. The Li-6 + n reaction requires neutrons with energies is the low MeV range or below. The Li-7 + n reaction is only significant when the energies are above 4 MeV.

2.2.2 Basic Principles of Fusion Weapon Design

2.2.2.1 Designs Using the Deuterium+Tritium Reaction

At ordinary material densities (e.g. liquid hydrogen, compressed hydrogen gas, or hydrogen-rich solids) reaction 1, the D+T reaction, is the only reaction that can occur to a significant extent at the temperature of an atomic bomb (50-100 million degrees, the temperature at the center of the Sun is only 14 million degrees!). This reaction can be used by simply allowing the fission explosion to heat up the fuel to fusion temperatures, requiring at most modest compression by the fission reaction itself. The high energy neutrons produced by this reaction are thus used to "boost" fission bombs.

The shortcoming with using the D+T reaction is that tritium is radioactive and decays at a rate of 5.5% year. This means that it does not occur in nature and must be manufactured through the use of reaction 5 in a nuclear reactor. In a weapons program tritium must compete with plutonium production in weapons production reactors, and it costs 80 times as much to produce one gram tritium compared with one gram of plutonium due its low atomic weight. It would be far too expensive to produce enough to make high yield fusion bombs. The decay also means that

it must be continually replenished. This reaction, using reactor manufactured tritium, has been used in low-yield neutron bombs where large amounts of tritium are not necessary

The production of tritium through reaction 5 can also be carried out in an atomic bomb, using the neutrons that escape from the critical assembly. This approach was used in the first hydrogen bomb tested by the Soviets. A large bomb cannot be manufactured using this method though because there are insufficient neutrons produced. On average each fission produces roughly one spare neutron, and releases 180 MeV of energy. If the spare neutron is captured by Li-6, producing one atom of tritium, which then fuses, we get a total energy production of 22.4 MeV. We would expect then that the fusion yield would be no more than about 10% of that of the fission trigger. If the fusion energy were the only contribution to the bomb's yield, then there would be no point in using this technique. The 14.1 MeV neutron from the D+T reaction can however cause fission in U-238, which is used in the fission tamper. This extra fission reaction doubles the bombs yield. Due to the low fusion yield inherent to the design this could be considered a type of fusion-boosted fission weapon.

2.2.2.2 Designs Using Other Fuels

It is militarily desirable to use fuels that are cheaper, and more stable than tritium. Deuterium, the sole fuel in reactions 2 and 3, is relatively cheap (especially considering its enormous energy content) and is completely stable. Pure deuterium has been used in at least one fusion weapon test - Ivy Mike, arguably the first true fusion weapon explosion in history (1 November 1952). Unfortunately deuterium, like all elemental hydrogen, is difficult to store. It must either be highly compressed, or liquified at extremely low temperatures. This problem can be overcome by combining the deuterium chemically with lithium to form lithium deuteride, a stable solid. An additional benefit is that through reactions 5 and 6, the lithium can itself participate in the fusion reaction.

To make use of these fuels, the slower reaction rates must be offset by compressing them to densities hundreds or thousands of times greater than those of normal conditions. At any given temperature the reaction rate goes up with the square of the density, a thousand-fold compression gives a million-fold reaction rate increase.

The work required to compress a gas is proportional to its temperature (at these pressures the physical strength of materials is negligible, and everything can be considered a gas). To minimize the work required for compression, or alternatively to achieve maximum compression for a given amount of work, it is important to keep the fusion fuel from getting hot until after the desired density is reached.

The key to making large fusion bombs is finding a way for using the energy of an atomic bomb trigger to compress a mass of deuterium sufficiently for the D-D reactions to become practical, followed by heating of the mass to ignition temperatures after the proper density has been achieved. The technique for doing this is staged radiation implosion, also called the Teller-Ulam configuration after its original joint inventors, Stanislaw Ulam and Edward Teller (also reinvented

independently by Andrei Sakharov and his associates, and by others in Britain, France, and China).

The Teller-Ulam configuration makes use of the fact that at the high temperatures of a fission bomb 80% or more of the energy exists as soft X-rays, not kinetic energy. The transport of energy by radiation from the fission core greatly exceeds the core's expansion rate (a mere 1000 km/sec or so). It is possible then to use this energy to compress, and ignite a physically separate mass of fusion fuel (the second stage) through radiation implosion before the expanding trigger disrupts it.

The principles of the Teller-Ulam configuration are more easily explained with the help of the diagram below. The bomb casing is roughly cylindrical, with the fission trigger at one end. The fusion fuel (lithium deuteride in the diagram) is a cylinder or ellipsoid wrapped in a pusher/tamper - a layer of very dense material (uranium or tungsten). Running down the axis of the fuel cylinder is a Pu-239 or U-235 rod, 2-3 cm or so in diameter. Lining the casing is a layer of plastic or plastic foam. Separating the trigger from the fuel package is a thick plug of dense material (again U or W).

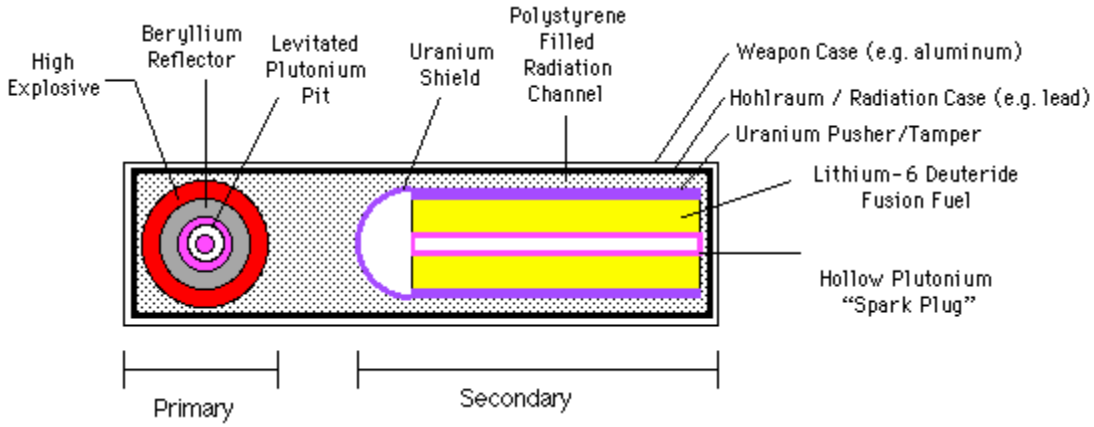

Components of the Teller-Ulam design:

- External Casing (made of structural material: steel, aluminum, plastic, etc.)
- Primary (fission trigger)
- Radiation Shield (high-Z material: uranium or tungsten; this may also contain boron-10 as a neutron absorber)
- Hohlraum or Radiation Case (high-Z material: uranium, lead, or tungsten, etc.)
- Radiation Channel (gap between the casing and the fusion pusher tamper; basically empty, often filled with plastic foam)
- Fusion Pusher/Tamper (high-Z material: natural/depleted uranium, HEU, tungsten, lead, etc.)
- Fusion Fuel (usually Li-6 deuteride; also natural lithium deuteride, liquid deuerium, etc.)
- Spark Plug (fissionable rod of HEU or plutonium)

When the trigger explodes, the X-rays escaping from the fission trigger fill the radiation channel, the space between the bomb casing and the fusion capsule, with a photon gas. This space is filled with plastic foam, essentially just carbon and hydrogen, which becomes completely ionized and transparent as the x-rays penetrate. The inner casing and outer capsule surfaces are heated to very high temperatures. The uranium shield between the trigger and the fusion capsule, and capsule pusher/tamper, prevents the fusion fuel from becoming heated prematurely.

Thermal equilibrium is established extremely rapidly, so that the temperature and energy density is uniform throughout the radiation channel. As the surface of the tamper becomes heated, it expands and ablates (blows off the fuel capsule surface). This ablation process, essentially a rocket turned inside out, generates tremendous pressure on the fuel capsule and causes an accelerating implosion. Thermal equilibrium assures that the implosion pressure is very uniformly distributed. The transparent carbon-hydrogen plasma retards the early expansion of the tamper and casing plasmas, keeping the radiation channel from being blocked by these opaque high-Z materials until equilibrium is fully established.

The force that compresses and accelerates the fusion fuel inward is provided solely by the ablation pressure. The other two possible sources of pressure - plasma pressure (pressure generated by the thermal motion of the plasma confined between the casing and the fuel capsule) and radiation pressure (pressure generated by thermal X-ray photons) do not directly influence the process.

The pressure exerted by the plasma causes cylindrical (or spherical) implosion of the fusion capsule, consisting of the pusher/tamper, fuel, and the axial fissionable rod. The capsule is compressed to perhaps 1/30 of its original diameter for cylindrical compression (1/10 for spherical compression), and thus reaches or exceeds 1000 times its original density. It is noteworthy that at this point the explosive force released by the trigger, an amount of energy sufficient to destroy a small city, is being used simply to squeeze several kilograms of fuel!

It is unlikely that the fissionable rod reaches such extreme compression however. Located at the center, it will experience an extremely violent shock wave that will heat it to high temperatures but compress it only modestly, increasing its density by a factor of 4 or so. This is sufficient to make the rod super-critical. Depending on the degree of symmetry, and the physics of the particular capsule collapse process higher densities are possible. Thermalized neutrons trapped in the fusion fuel, which are left over from the intense fission neutron flux, initiate a chain reaction as soon as the rod becomes critical. The rod fissions at an accelerating rate as it, and the rest of the fuel capsule continue to implode and acts as the fusion "spark plug". Combined with the high temperatures generated by the convergent shock wave, this raises the temperature of the fusion fuel around the rod high enough to initiate the fusion reaction. Self-supporting fusion burning then spreads outward. The fusion tamper prevents the escape of thermal radiation from the fuel. As the temperature rises the fusion reactions accelerate, enhancing the burn efficiency considerably. The temperatures generated by fusion burning can exceed 300 million K, considerably more than that produced by fission.

The fuel in the fusion capsule consists of lithium deuteride that may be enriched in the Li-6 isotope (which makes up 7.5% of natural lithium). Natural lithium has been used with success in fusion bomb designs, but modern light weight designs seem to use lithium enriched in Li-6.

There is some tritium generated by the fission neutrons, but as noted above the contribution to bomb yield is insignificant. Far more tritium is produced by the D+D reactions, either directly by reaction 3, or by reaction 5 via the neutrons produced in reaction 2.

Since the D+T reaction rate is so high, and there is large excess of deuterium, the tritium is consumed almost as fast it is produced. The 14.1 MeV neutrons can also produce large amounts of tritium from Li-7 through reaction 6.

A large part of the fusion fuel can be burned before expansion quenches the reaction by reducing the density, which takes some 20-40 nanoseconds. The power output of a fusion capsule is noteworthy. The largest bomb ever exploded had a yield of 50 Mt, almost all produced by its final fusion stage. Since 50 Mt is 2.1x10^17 joules, the power produced during the burn was around 5.3x10^24 watts. This is more than one percent of the entire power output of the Sun (4.3x10^26 watts)!! The peak output was possibly even greater.

The 2.45 MeV and 14.1 MeV neutrons that escape from the fusion fuel can also contribute greatly to bomb yield by inducing fission in the highly compressed fusion tamper. This extra boost can release most of the explosion energy, and commonly accounts for half of the yield of large fission-fusion-fission bombs and can reach at least 85% of the total yield, although in most designs it is more like 50%.

The Teller-Ulam fusion bomb described so far is called a "two stage bomb". The fission trigger (the first stage) compresses the fusion capsule (the second stage). As powerful as the trigger is, there is a limit to how large a capsule it can compress in the brief time available. If a still bigger bomb is desired, then the explosion of the fusion secondary can be used to compress and explode a larger third stage. Each stage can be 10-100 times the size of the previous stage. The 50 Mt bomb mentioned above was a three stage weapon.

3.0 Matter, Energy, and Radiation Hydrodynamics

This section provides background in the fundamental physical phenomena that govern the design of nuclear weapons, especially thermonuclear weapons. This section does not address nuclear physics which are introduced in Section 2, and discussed further in Section 4. It addresses instead the behavior of matter at high densities and temperatures, and the laws controlling its flow.

Although the reader may be able to follow the discussions of physics and design in Section 4 without it, familiarity with the principles discussed here is essential for genuine insight into the

design of nuclear weapons. Since the same principles tend to crop up repeatedly in different contexts, and it is inconvenient to explain basic physics while addressing engineering considerations, I provide an overview of the non-nuclear physical principles involved below. Section 2 provides a discussion of the nuclear reactions involved.

Readers with a grounding in physics will find much of this discussion too elementary to be of interest. Please skip anything you are already familiar with.

- 3.1 Thermodynamics and the Properties of Gases
- 3.2 Properties of Matter
- 3.3 Interaction of Radiation and Matter
- 3.4 Hydrodynamics
- 3.5 Radiation Hydrodynamics
- 3.6 Shock Waves in Non-Uniform Systems
- 3.7 Principles of Implosion
- 3.8 Instability

3.1 Thermodynamics and the Properties of Gases

Thermodynamics concerns itself with the statistical behavior of large collection of particles, a substantial quantity of matter for example. A literal reading of the term "thermodynamics" implies that the topic of discussion is the motion of heat. In fact, thermodynamics specifically addresses the condition of thermal equilibrium, where the motion of heat has ceased. The principle motion of interest is the randomized motion of the particles themselves, which gives rise to the phenomenon called "heat". It is this equilibrium condition of random motion that can be accurately characterized using statistical techniques. It is important to realize that the topic of thermodynamics is the study of the properties of matter at least as much as it is of heat.

There are two very important principles of thermodynamics, called the First and Second Laws of Thermodynamics. These are often stated in the form:

1. The total energy in the universe is constant;
2. The entropy of the universe is always increasing.

A more useful statement of the First Law in practical situations is to say that the change in total energy of a system is equal to the work done on the system, plus the heat added to the system. The Second Law states that the amount of heat in a closed system never decreases. The implications of these laws are discussed further below.

3.1.1 Kinetic Theory of Gases

The gaseous state is the simplest form of matter to analyze. This is fortunate, since under the extreme conditions encountered in chemical and nuclear explosions, matter can usually be treated as a gas regardless of its density or original state.

The basic properties of a gas can be deduced from considering the motions of its constituent particles (the kinetic theory). The pressure exerted by a gas on a surface is caused by the individual molecules or atoms bouncing elastically off that surface. This pressure is equal to the number of molecules striking the surface per unit of time, multiplied by the average momentum of each molecule normal (i.e. at right angles) to the surface. The number of impacts per second is proportional (~) to the particle density of the gas (rho, particles per unit volume), and how fast the molecules are traveling (the average molecular velocity v):

Eq. 3.1.1-1

$$P \propto \rho \bullet v \bullet momentum$$

The average momentum is proportional to v times the mass of the particles (m). The pressure is thus:

Eq. 3.1.1-2

P ~ rho*v*v*m.

Actually we can state that:

Eq. 3.1.1-3

P = rho*v*v*m/3

since the contribution of molecular velocity normal to the surface in three dimensions is 1/3 of the squared magnitude.

Since v*v*m/2 is the particle kinetic energy (KE_p), we can also say:

Eq. 3.1.1-4

P = rho*KE_p*(2/3)

That is, pressure is proportional to the average particle kinetic energy and the particle density, or equal to two-thirds of the total kinetic energy, KE, in a given volume of gas (the kinetic energy density). This is usually expressed as:

Eq. 3.1.1-5

P = 2/3(KE/V), or

PV = 2/3 KE,

where P is the pressure.

Now the thing we call *temperature* is simply the average kinetic energy of the particles of a gas. A constant of proportionality is used to convert kinetic energy, measured in joules or ergs, into degrees Kelvin (K). Together these considerations give us the *Ideal Gas Law*:

Eq. 3.1.1-6

PV = NkT, where

P = pressure, V = volume, N = number of particles, k = Boltzmann's constant (1.380 x 10^-16 erg/degree K), and T = temperature. N/V is of course the particle density (designated n).

The constant factor 2/3 was absorbed by Boltzmann's constant. As a result, if we want to express the average particle kinetic energy of a gas at temperature T we must say:

Eq. 3.1.1-7

$KE_p = 3/2\ kT$

An *ideal gas* (also called a *perfect* gas) is one in which there are no interactions (that is, repulsive or attractive forces) between atoms. For such a gas, the Ideal Gas Law holds true. The simplest case of a perfect gas is a *perfect monatomic* gas, one in which all of the energy in the gas is in the form of particle motion (i.e. the particles themselves do not absorb any energy). This is the only case we have considered so far. Helium or argon are examples of ideal monatomic gases to a very good approximation (they are monatomic, and attractive forces only become significant close to their liquefaction temperatures).

Molecular or polyatomic gases, ones in which the particles are molecules of two or more atoms, can absorb energy through rotation and vibration. Such gases are not monatomic, but they are still ideal. Under some conditions gases can absorb energy internally by other processes, like ionization, which violate ideal gas behavior. When conditions are such that attractive forces become significant (near liquid or solid condensation points) the ideal gas law also breaks down.

Perfect monatomic gases are of special interest to us here, not only because they are particularly simple to analyze, but because under many extreme physical regimes all matter tends to behave like a perfect monatomic gas (kinetic energy dominates other forms of energy present).

3.1.2 Heat, Entropy, and Adiabatic Compression

Simply put, heat is the random motion of the particles in matter. In common usage we talk about something with a higher temperature as being "hotter". However temperature is *not* a universal measure of the thing we call heat. Suppose we take a container of a perfect gas, and we squeeze it and reduce its volume. To squeeze it and compress the gas we must do work which, by the First Law of Thermodynamics, is added to the internal energy of the gas. Since this is a perfect gas, all of the added energy appears as kinetic energy. That is, the temperature goes up. But have we actually added heat to make it hotter?

The answer is no. We can get the energy back in the form of work, by letting it expand back to its original volume. The temperature will also drop back to the original state. This compression process (called *adiabatic* compression) is *reversible* since we can return to the original state.

To increase the temperature of the container of gas without changing its volume, we must place it in contact with something that is hotter. The heat diffuses from the hotter object to the container. As the gas in the container warms, the hotter object grows cooler.

How can we return the gas to its original state? We must place it in contact with something that is colder than the original gas temperature. The heat then diffuses to the colder object. Although the *gas in the container* is now in its original state, the whole system is not. The hotter object is cooler, the colder object is warmer. This process is *irreversible* (we say "entropy of the system has increased").

Temperature is a measure of heat in a gas only at constant volume. The generalized measure of heat is *entropy*. Entropy is defined as the ratio of the total energy of a system to its temperature. As heat is added to a system this ratio increases. Work done on the system leaves the ratio unchanged.

Adiabatic compression is compression where the entropy is constant (no heat is added or removed). If flows of heat occur, then the process is non-adiabatic and causes irreversible change.

3.1.3 Thermodynamic Equilibrium and Equipartition

I have just talked about heat flowing from hotter objects to colder ones. This process implies that a system of objects tends to move to a state where all of the objects are at the same temperature. When this occurs, heat ceases to flow. Such a state is called "thermodynamic equilibrium", and all systems tend to evolve toward this equilibrium naturally. The faster heat can flow in the system, the faster this equilibrium is reached.

The idea of thermodynamic equilibrium is extremely general. It applies not only to "objects" - physically separate parts of a system - but *all* parts of a system - separate or not.

For example in a mixture of particles of different types, different gas molecules say, each type of particle will be in equilibrium with the others. That is, they will have the same temperature - the same average kinetic energy. If each type of particle has a different mass from the others, then each must also have a unique average velocity for the kinetic energies of each type to be equal. One implication of this is that when a gas becomes ionized, the electrons knocked loose become separate particles and will come into thermodynamic equilibrium with the ions and un-ionized atoms. Since they are much lighter than atoms or ions, their velocities will be much higher.

We have also already applied the equilibrium principle in deriving the Ideal Gas Law. The total kinetic energy was divided equally among the three spatial directions of motion, e.g. they were in equilibrium with each other. These spatial directions are called the "degrees of freedom" of a monatomic perfect gas. Since the kinetic energy of a particle in such a gas is 3kT/2, each degree of freedom accounts for kT/2 energy per particle. This is also true of polyatomic gases, which have additional degrees of freedom (e.g. from vibration and rotation). Each available degree of freedom will have kT/2 energy when in equilibrium. This is the theorem of *equipartition of energy*.

The actual number of available degrees of freedom in a polyatomic gas may vary significantly with temperature due to quantum-mechanical considerations. Each degree of freedom has a

characteristic energy of excitation, and if the value of kT/2 is not large enough then the excitation of a given state will be negligible.

3.1.4 Relaxation

To reach equilibrium between different particles and different degrees of freedom in a system, the different parts of the system must be able to exchange energy. The rate of energy exchange determines how long it takes to establish equilibrium. The length of this equilibrating period is called the *relaxation time* of the system. A complex system will typically have several relaxation times for different system components.

The farther a degree of freedom is from equilibrium, the faster it will converge toward the equilibrium state. Conversely, as it approaches equilibrium, the rate of convergence declines. This is expressed by the standard relaxation equation:

Eq. 3.1.4-1

$$dE/dt = (E_{eq} - E)/t_{relax}$$

where E is the measure of the current energy of the degree of freedom (avg. kinetic energy, temperature, number of particles excited, etc.), E_eq is the equilibrium value, and t_relax is the relaxation time.

The solution of this linear differential equation shows us that the difference between the current state and the equilibrium state declines exponentially with time:

Eq. 3.1.4-2

$$E = E_{init}*Exp[-t/t_{relax}] + E_{eq}*(1 - Exp[-t/t_{relax}])$$

Over each time interval t_relax, the difference E - E_eq declines by a factor of 1/e. Although according to this equation complete equilibrium is never formally reached, over a finite (usually small) number of relaxation intervals the difference from equilibrium becomes undetectable.

What determines the value of t_relax? This is determined by how frequently a member of a degree of freedom can be expected to undergo an energy exchange event, and how effective that event is in transferring energy.

For particles of similar mass, a single collision can transfer essentially all of the kinetic energy from one particle to the other. The relaxation time for bringing two populations of particles with different kinetic energies into equilibrium is thus the average time between collisions. In air at normal temperatures and pressures, this time is about 0.1 nanoseconds. At higher densities and temperatures, the distances traveled between collisions is shorter, and the velocities are higher, so the time is correspondingly shorter.

If colliding particles have greatly different masses, then the efficiency of each collision in exchanging energy is reduced by a factor equal to the mass ratio. In the case of electrons and ions, since electrons are lighter than nucleons by a factor of 1836 (about) this ratio is 1/(1836*A), where A is the atomic mass. Unless the temperature of the electrons is much colder than that of the ions though, the actual relative relaxation rate is much higher than this would indicate because of the high velocities of the light electrons. If they are not too far from equilibrium, the actual relaxation time ratio between electrons and ions, and ions alone is about equal to the square root of the mass ratio: 1/(1836*A)^0.5.

3.1.5 The Maxwell-Boltzmann Distribution Law

So far we have talked about the *average* velocity and kinetic energy of a particle. In reality, no particle will have exactly the average energy. Even if we created a system in which every particle initially had exactly the same energy (*all* were average), within a single relaxation interval the energy would be dramatically redistributed. Within a few more intervals a stable continuous energy distribution would be established.

Statistical mechanics shows that the actual equilibrium distribution of particle energies can be described by the distribution law worked out first by Maxwell and refined by Boltzmann. The function creates a roughly bell-shaped curve, with the peak (most probable) energy at kT. The function declines exponentially away from the peak, but never (formally) goes to zero at any energy greater than zero, so small numbers of both very fast and very slow particles are present in an equilibrium gas.

The Maxwell-Boltzmann distribution for energy is:

Eq. 3.1.5-1

$$dN/dE = N*2*Pi*(1/(Pi*kT))^(3/2)\ Exp(-E/kT)\ E^(1/2)$$

where N is the number of particles present. Integrating the above equation over a given energy range gives the number of particles in that range.

Most of the terms in the above equation are simply normalizing factors to make the integral come out right (after all, integrating from zero energy to infinity must equal N). The factor that actually determines the distribution law is called the Boltzmann factor: Exp(-E/kT). This distribution factor applies to any system of particles where each energy state is equally important in a statistical sense (that is, no statistical weight is applied to any energy state). A gas where this is true (like the gases treated by classical kinetic theory) can be called a Boltzmann gas. There are two other types of gases that follow different distribution laws which will be discussed later - the Bose gas and the Fermi gas.

Below is a plot for the Maxwell-Boltzmann particle density distribution, dN, with kT=1. The peak value of dN is at particle energy kT/2, but since the energy density distribution is proportional to dN*E, the peak of the energy density distribution is actually 3kT/2.

Figure 3.1.5-1. Maxwell-Boltzmann Particle Density Distribution for Energy

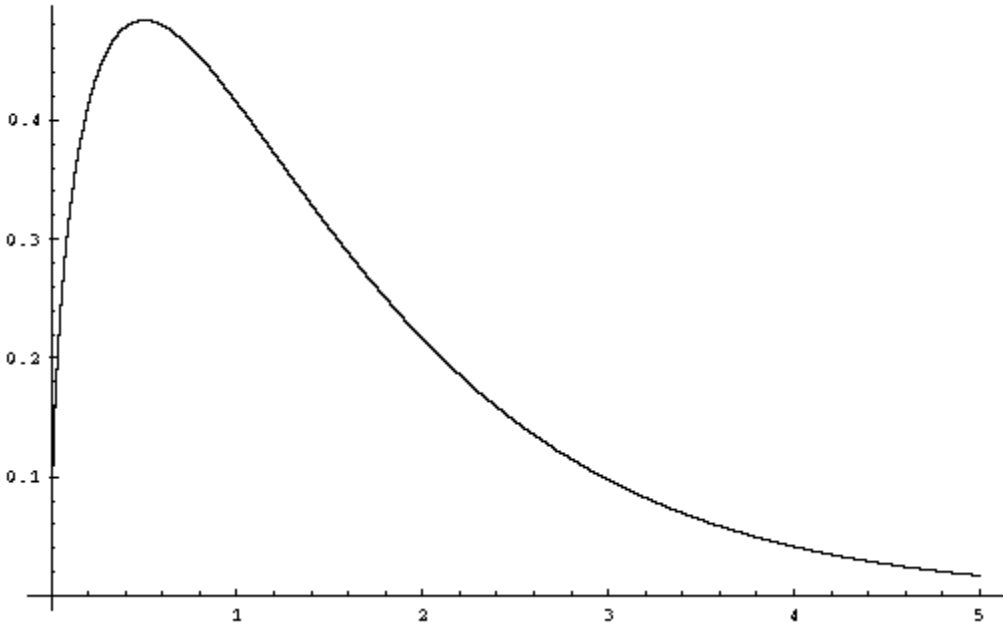

3.1.6 Specific Heats and the Thermodynamic Exponent

The Ideal Gas Law describes the properties of gases with respect to temperature, that is the kinetic energy of motion. How do we describe the properties of a gas with respect to the *total internal energy*? In the case of a monatomic gas this is easy of course, since the kinetic energy is the total internal energy:

$$PV = NkT = 2/3\ KE$$

How should we handle polyatomic gases? Before we can do this we need some convenient way of measuring the thermodynamic properties of different gases.

As I have explained above, an ideal gas with additional degrees of freedom has a larger internal energy than does an monatomic gas at the same temperature. This internal energy, designated U, is also proportional to the (absolute) temperature (this is another common way of expressing the concept of *ideal gas*). This allows us to establish a constant for each gas that describes how much thermal energy is required to raise its temperature a fixed amount. This constant is called the *specific heat*.

There are actually two commonly used specific heat definitions for gases, the specific heat at constant volume (c_v) and the specific heat at constant pressure (c_p). C_v measures the amount of energy required to raise the temperature in a sealed, fixed volume container. In such a container heating also causes the pressure to rise. C_p measures the amount of energy required to raise the temperature of a gas that is allowed to expand sufficiently to maintain constant pressure.

These two specific heats are not independent. In fact, the ratio between them is fixed by the number of degrees of freedom of the gas. This gives us the constant that we use for describing the thermodynamic properties of a gas - the thermodynamic exponent. This constant is represented by the lower case Greek letter gamma. It is defined by:

Eq. 3.1.6-1

$$\gamma = c_p/c_v$$

and is equal to 5/3 for a monatomic gas. The thermodynamic exponent has many other names such as the adiabatic index, adiabatic exponent, isentropic exponent, and the polytropic exponent.

Recall that the internal energy density of a monatomic gas (KE/V) is given by:

$$P = 2/3\ (KE/V)$$

Since KE is also the total internal energy we can say:

Eq. 3.1.6-2

$$P = 2/3\ U/V$$

The factor 2/3 happens to be equal to gamma minus 1: 5/3 - 1 = 2/3. This is a special case of a law valid for all perfect gases. We can thus write the general law:

Eq. 3.1.6-3

$$P = (\text{gamma} - 1)*(U/V)$$

Why is gamma called the thermodynamic or adiabatic *exponent*? It is because of the following relationship that describes the state of matter undergoing adiabatic compression:

Eq. 3.1.6-4

$$P(V^{\text{gamma}}) = \text{constant}$$

The constant is determined by the gases' original entropy. This is sometimes called the polytropic law.

The thermodynamic exponent determines the compressibility of a gas. The larger the value of gamma, the more work is required to reduce the volume through adiabatic compression (and the larger the increase in internal energy). An infinitely compressible gas would have an exponent of 1.

The thermodynamic literature often uses P-V diagrams that plot pressure (P) versus volume (V). A plot of the adiabatic function on a P-V diagram produces an adiabatic curve, also called an

isentropic curve since every point on the curve has the same entropy. In contrast, isothermal curves lie below the adiabatic curve with increasing pressure (assuming they start at the same P-V state) since a gas must lose entropy to maintain the same temperature. Curves where entropy is increased with increasing pressure lie above (these are Hugoniot curves which will be discussed further in connection with shock waves).

Gamma for a gas is related to the number of degrees of freedom (n_f) by:

Eq. 3.1.6-5

$$\text{gamma} = (2/n_f) + 1$$

Thus a monatomic perfect gas is 2/3 + 1 = 5/3 as noted above. A diatomic gas has a maximum of 7 degrees of freedom, but only some of them may be excited at a given temperature, with more states being excited at higher temperatures.

If a gas *releases* energy during compression, thus adding additional kinetic energy, (due to a chemical reaction for example) then it will have a higher value of gamma.

Some example values of gamma are given in the table below.

Table 3.1.6-1. Examples of Thermodynamic Exponents for Gases

Material	Exact Value	Approx. Value
Detonating Explosive Gas Mixture	-	2.5-3.0
Perfect Monatomic Gas	5/3	1.667
Air	7/5	1.400
Photon Gas	4/3	1.333

Diatomic Gas (fully excited)	9/7	1.286
Infinitely Compressible Gas	1/1	1.000

3.1.7 Properties of Blackbody Radiation

The equipartition of energy in an equilibrium system also extends to radiant energy present in the system. Photons are emitted and absorbed continually by matter, creating an equilibrium photon gas that permeates it. This photon must have the same temperature as the rest of the system also.

The energy distribution in an equilibrium photon gas is determined by quantum mechanical principles known as Bose-Einstein statistics. Photons belong to a class of particles called bosons that, by definition, obey these statistics. A key feature of bosons is that they prefer to be in the same energy state as other bosons. A photon gas is thus an example of a Bose gas. The distribution factor for Bose-Einstein statistics is: 1/(Exp(E/kT) - 1).

This fact gives rise to an energy distribution among the particles in a photon gas called the blackbody spectrum which has a temperature dependent peak reminiscent of the Maxwell-Boltzmann distribution. The term "blackbody" refers to the analytical model used to derive the spectrum mathematically which assumes the existence of a perfect photon absorber or (equivalently) a leakless container of energy (called in German a "hohlraum").

The kinetic theory of gases can be applied to a photon gas just as easily as it can to a gas of any other particle, but we need to make a few adjustments. From Eq. 3.1.1-1 we had:

P ~= rho*v*momentum

Which gave us Eq. 3.1.1-3:

P = rho*v*v*m/3

once we had substituted m*v to represent the momentum of a particle. Since photons have zero mass, we must use a different expression to express the momentum of a photon. This is given by:

Eq. 3.1.7-1

$$momentum_{photon} = E_{photon}/c$$

where E_photon is the photon energy, and c is the photon velocity (i.e. the speed of light, 2.997 x 10^10 cm/sec). It is interesting to compare this to an equivalent expression for massive particles: momentum = 2*KE/v. Substituting Eq. 3.1.7-1, and the photon velocity, into Eq. 3.1.1-3 give us:

Eq. 3.1.7-2

$$P_{rad} = \rho * c * (E_{photon}/c)/3 = \rho * E_{photon}/3$$

Since rho*E_photon is simply the energy density of the photon gas, we can say:

Eq. 3.1.7-3

$$P_{rad} = (U_{rad}/V)/3$$

From Eq. 3.1.6-3 it is clear that:

Eq. 3.1.7-4

$$\gamma_{rad} = 1 + 1/3 = 4/3$$

We can relate the energy density of a blackbody to the thermal radiation emissions (energy flux) from its surface (or from a window into an energy container). Assuming the energy field is isotropic, the flux is simply the product of the energy density and the average velocity with which the photons emerge from the radiating surface. Of course all of the photons have a total velocity equal to c, but only photons emitted normal to the surface (at right angles to it) emerge at this velocity. In general, the effective velocity of escape is Cos(theta)*c, where theta is the angle between the light ray and the normal vector. Now the fraction of a hemisphere represented by a narrow band with width d_theta around theta is Sin(theta)*d_theta. Integrating this from zero to 90 degrees gives (in Mathematica notation):
(U_rad/V)*Integrate[c*Cos(theta)*Sin(theta),{theta,0,90}] = (U_rad/V)*c/2 Since the flux is isotropic, half of it is flowing away from the surface. We are only concerned with a the flux flowing out so we must divide it by another factor of two. This gives:

Eq. 3.1.7-5

$$S = c*(U_{rad}/V)/4$$

where S is the flux (emission per unit area).

At equilibrium the radiation energy density is determined only by temperature, we want then to have a way for relating a temperature T to U_rad/V. Using Eq. 3.1.7-3 and a mathematically precise statement of the second law of thermodynamics, it is easy to show that U_rad/V is proportional to T^4. The standard constant of proportionality, called the Stefan-Boltzmann constant and designated sigma, is defined so that:

Eq. 3.1.7-6

$$U_{rad}/V = (4*\sigma/c)*T^4$$

This is a convenient way of formulating the constant, because it allows us to say:

Eq. 3.1.7-7

S = sigma*T^4

Eq. 3.1.7-7 is known as the Stefan-Boltzmann Law. The Stefan-Boltzmann constant is derived from Planck's constant and the speed of light. It has the value 5.669 x 10^-5 erg/sec-cm^2-K, with T in degrees K. Equation 3.1.6-3 of course becomes:

Eq. 3.1.7-8

P_rad = ((4*sigma)/(3*c))*T^4

It can easily be seen from the Stefan-Boltzmann Law that the amount of radiant energy present varies dramatically with temperature. At room temperature it is insignificant, but it grows very rapidly. At sufficiently high temperatures, the energy present in the blackbody field exceeds all other forms of energy in a system (which is then said to be "radiation dominated"). The average photon energy is directly proportional to T, which implies the photon density varies as T^3. In radiation dominated matter we can expect the number of photons present to be larger than the number of all other particles combined.

If both ordinary particles and photons are present, we have a mixture of Boltzmann and Bose gases. Each contribute independently to the energy density and pressure of the gas mixture. Since the kinetic energy pressure for a perfect gas is:

PV = NkT -> P = nkT

and the kinetic energy is:

Eq. 3.1.67-9

P = 2/3 KE/V = 2/3 U_kin/V -> U_kin/V = 3/2 nkT

we have:

Eq. 3.1.7-10

$$P_{total} = nkT + ((4*\sigma)/(3*c))*T^4$$

and:

Eq. 3.1.7-11

$$U_{total}/V = 3/2\ nkT + ((4*\sigma)/c)*T^4$$

We can calculate the relative kinetic and radiation contributions to both pressure and energy at different particle densities and temperatures. For example in hydrogen at its normal liquid density, radiation energy density is equal to the kinetic energy of the ionized gas at 1.3 x 10^7 degrees K.

The energy distribution of the radiation field with photon energy is given by Planck's Law, which is usually stated in terms of photon frequency instead of energy. The energy of a photon of frequency nu is simply:

Eq. 3.1.7-12

$$E_{phot} = h*\nu$$

where nu is in hertz (c/wavelength), and h is Planck's constant (6.62608 x10^-27 erg-sec). Planck's law (aka the Planck function) is usually given as:

Eq. 3.1.7-13

$$dE/dnu = ((8*Pi*h*nu^3)/c^3) * (1/(Exp((h*nu)/kT) - 1))$$

where dE/dnu the energy density/frequency derivative. The last factor in the equation is of course the Bose-Einstein distribution factor. Integrating over a range of nu gives the energy density in that frequency range. For our purposes, it is often more convenient to express the energy density in terms of photon energy rather than frequency:

Eq. 3.1.7-1e

$$dE/dE_phot = ((8*Pi*E_phot^3)/(h^3\ c^3)) * (1/(Exp(E_phot/kT) - 1))$$

The Planck distribution always has its peak (the maximum spectral power) at h*nu_max = 2.822 kT, while 50% of the energy is carried by photons with energies greater than 3.505 kT, and 10% of the energy is above 6.555 kT. Most of the energy in the field is thus carried by photons with energies much higher than the average particle kinetic energy.

Below is a diagram of the Planck function at a temperature of 1 KeV, plotting the spectral energy density against the photon energy in KeV.

Figure 3.1.7-1. Blackbody Spectrum

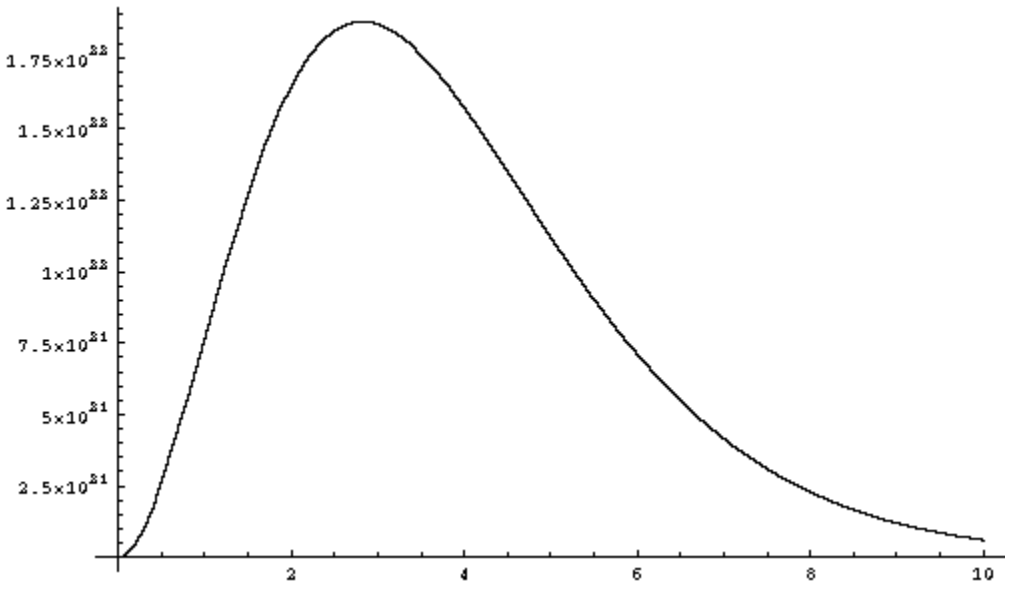

3.2 Properties of Matter

I have already discussed one state of matter - gases - at some length. In this section I shift to the application of thermodynamic principles to other states of matter, and discuss some properties that are not strictly thermodynamic in nature.

3.2.1 Equations of State (EOS)

An *equation of state* (EOS) provides a complete description of the thermodynamic properties of a substance; i.e. how the density, pressure, and internal energy of a substance relate to each other. The Ideal Gas Law is a special case of an equation of state for gases. The generalized gas EOS given previously:

P = (gamma - 1)*(U/V)

expands the ideal law to all gases if an appropriate value of gamma is chosen.

Such a simple law is not really adequate for real substances over widely varying conditions. Even with comparatively simple substances such as gases, the effective value of gamma can change. As molecular gases increase in temperature more degrees of freedom may become excited, the gases may disassociate into atoms, and the atoms may become ionized. All of these processes drive down the value of gamma by absorbing energy that would otherwise appear as kinetic motion. By considering the regime of interest, we can usually choose a suitable value of gamma to permit the use of the simple gas equation. More sophisticated

approaches are to provide terms for each mechanism that contributes to the total internal energy.

3.2.2 Condensed Matter

The term "condensed matter" refers to two of the three common states of matter: solids and liquids. It describes the fact that the matter is not gaseous, it has condensed to a compact form bound together by interatomic attractive forces. At zero pressure (or atmospheric pressure, which is the same thing for practical purposes) condensed matter exists in equilibrium. The negative pressure generated by the binding forces is exactly balanced by positive forces generated by the mutual repulsion of the outer electron shells (Coulomb repulsion) and the thermal motion of the atoms. Condensed matter thus does not expand to infinitely low density under zero pressure like a gas, it has a definite zero pressure density.

Another important difference between condensed matter and gases is the strength of the internal repulsive forces. Coulomb repulsion is much stronger than the kinetic forces produced by thermal motion in gases under ordinary conditions, which agrees with the common experience that a brick is harder to compress than air.

If the thermal energy in matter is low enough, the position of atoms are held in fixed positions by Coulomb repulsion - it is a solid. When the thermal vibrations become sufficiently energetic, the atoms break free from fixed lattice positions and can move around and the solid melts.

3.2.3 Matter Under Ordinary Conditions

The operative definition of "ordinary conditions" I am using here are the conditions under which condensed matter exists. It will be shown below that regardless of its composition or initial state, at sufficiently extreme conditions of pressure or temperature matter ceases to be condensed and tends to behave like a perfect gas.

There are standard definitions of ordinary conditions: STP or Standard Temperature and Pressure (0 degrees C temperature, 760 mm Hg pressure); or 0 degrees K and zero pressure. The conditions of normal human experience do not deviate much from STP, and the properties of most substances under these conditions are abundantly documented in numerous references. For our purposes "ordinary conditions" extends up to temperatures of a few tens of thousands of degrees C, and pressures in the order of several megabars (millions of atmospheres). For comparison the conditions in the detonation wave of a powerful high explosive do not exceed 4000 degrees C and 500 kilobars; the pressure at the center of the Earth is approximately 4 megabars.

Under our "ordinary conditions" the thermal energy of matter remains below both the binding and compressive energies. In this range matter is not appreciably ionized. Its mechanical strength is small compared to the pressures of interest, and can usually be neglected.

Since increasing pressure also strengthens the repulsive forces between atoms by forcing them closer together, the melting point goes up as well. In the megabar range matter remains solid

even at temperatures of 20-30,000 degrees C. However it usually does not matter whether the condensed state is liquid or solid, the energy absorbed in melting being too small to notice compared to the compressive and thermal energies.

Some materials undergo abrupt phase changes (discontinuous changes in structure and density) in this realm. When phase changes occur with escalating pressure, atoms suddenly rearrange themselves into lower energy configurations that are denser. For example iron undergoes a phase change at 130 kilobars. The transformation of delta phase plutonium alloys into the denser alpha phase at pressures of a few tens of kilobars is of particular significance.

Despite these differences and added complexities, we can still produce reasonable approximations for condensed matter equations of state using a "Gamma Law" similar to the gas law:

Eq. 3.2.3-1

$$P + P_0 = (\gamma - 1)*U*(\rho/\rho_0)$$

where P is the compression pressure at the state of interest, P_0 is the internal pressure at STP (or some other reference state), rho is the density of the state of interest, and ρ_0 is the reference density. Note that P_0 is exactly balanced by the negative binding pressure under reference conditions.

This gives us an adiabatic law for condensed matter:

Eq. 3.2.3-2

$$(P + P_0)*(\rho_0/\rho)^\gamma = \text{constant} = P_0$$

Another useful relationship is the equation for internal energy per unit mass (E) rather than energy per volume (U):

Eq. 3.2.3-3

E = U/rho_0 = (P + P_0) / ((gamma - 1)*rho)

The value of gamma, the "effective thermodynamic exponent", must be determined from experimental data. Unfortunately the value of gamma is not constant for condensed matter, it declines with increasing density and pressure. It is virtually constant below 100 kilobars, but the decline is in the range of 15-30% at 2 megabars. Although the rate of decline varies with substance, the low pressure value still gives a reasonable indication of compressibility of a substance at multi-megabar pressures. A common assumption in high pressure shock work is that the product of (gamma-1) and density is constant:

Eq. 3.2.3-4

(gamma_0 - 1)*rho_0 = (gamma - 1)*rho

an approximation which seems to work fairly well in practice. Using this approximation, Eq. 3.2.3-3 becomes particularly convenient since E varies only with P.

Eq. 3.2.3-5

E = (P + P_0) / G

where

G = (gamma_0 - 1)*rho_0 = constant

The thermodynamic exponent is usually represented in the literature of solid state or condensed matter physics by the "Gruneisen coefficient" designated with the upper case Greek letter GAMMA. The relationship between them is:

Eq. 3.2.3-6

GAMMA = gamma - 1

Representative values of P_0, density_0, and low pressure gamma for some materials of particular interest are given below:

Table 3.2.3-1. Gamma-Law Equations of State for Selected Materials

Material	Density_0	Gamma	P_0 (kilobars)
Gold	19.24	4.05	510
Aluminum	2.785	3.13	315
Copper	8.90	3.04	575
Detonating High Explosive	-	3.0	-
Uranium	18.90	2.90	547
Tungsten	19.17	2.54	1440
Beryllium	2.865	2.17	604
Lithium	0.534	2.1	-
Zirconium	6.49	1.771	580
Perfect Monatomic Gas	-	1.667	-

3.2.4 Matter At High Pressures

As pressures continue to increase above several megabars, the electronic structure of the atom begins to break down. The Coulomb forces become so strong that the outer electrons are displaced from the atomic nuclei. The material begins to resemble individual atomic nuclei swimming in a sea of free electrons, which is called an electron gas. This gas is governed by quantum mechanical laws, and since electrons belong to a class of particles called fermions (which obey Fermi-Dirac statistical laws), it is an example of a Fermi gas.

In contrast to the Bose-Einstein gas of photons, where particles prefer to be in the same energy state, fermions cannot be in the same energy state. Even at absolute zero, the particles in a Fermi gas must have non-zero energy. The distribution factor for Fermi statistics is: 1/(Exp(E/kT) + 1).

If all of the electrons are in their lowest energy state, which means the gas is cold (no additional thermal energy), it is said to be Fermi degenerate. A fully degenerate state is the lowest energy state that a Fermi gas can be in. A degenerate Fermi gas is characterized by the Fermi energy, the highest energy state in the gas. This is given by:

Eq. 3.2.4-1

$$E_{Fermi} = 5.84 \times 10^{-27} (n^{(2/3)}) \text{ erg} = 3.65 \times 10^{-15} n^{(2/3)} \text{ eV}$$

where n is the electron density (electrons/cm^3). The average electron energy is:

Eq. 3.2.4-2

$$E_{Favg} = 3/5 E_{Fermi}$$

and the pressure produced, the Fermi pressure, is:

Eq. 3.2.4-3

$$P_{Fermi} = 2/3 n*E_{Favg} = 2/5 n*E_{Fermi} = 2.34 \times 10^{-33} (n^{(5/3)}) \text{ bars}$$

Note that relationship between the average energy and the pressure is precisely the same as that for a classical perfect gas.

When the average electron energy exceeds the binding energy of electrons in atoms, then the electrons behave as a Fermi gas. If only some of the outer electrons are loosely bound enough meet this criterion, then only these electrons count in determining the electron density in the equations above, the remainder continue to be bound to the atomic nuclei.

The Fermi energy of a gas is sometimes characterized by the "Fermi Temperature" (or degeneracy temperature). This is defined as T_Fermi such that:

Eq. 3.2.4-4

$$kT_{Fermi} = E_{Fermi}$$

This is not the actual temperature of the gas. Its significance is that if the kinetic temperature is substantially lower than T_Fermi then the kinetic energy is small compared to the Fermi energy and the gas can be treated reasonably well as if it were completely degenerate ("cold").

To illustrate these ideas here are some examples: Uranium at twice normal density (37.8 g/cm^3) would have a Fermi energy of 156 eV, and a pressure of 895 megabars. This is much higher than the real pressure required to achieve this density (5.0 megabars), and indicates that the uranium is not a Fermi gas at this pressure.

A pressure of 100,000 megabars corresponds to a Fermi energy of 1034 eV, and an average energy of 621 eV. The average energy is about the same as the ionization energy for uranium's 38th ionization state. Thus we can expect about 41% of uranium's electrons to dissociate at this pressure, and contribute to the electron gas density (1.5 x 10^26 electrons/cm^3). This gives a density estimate of 1560 g/cm^3.

Deuterium at 1000 times normal liquid density (159 g/cm^3) is a true Fermi gas. It has E_Fermi = 447 eV (T_Fermi = 5.2 million degrees K), far higher than its ionization energy (13.6 eV), and P_Fermi = 12,500 megabars. What this says is that at least 12.5 gigabars of pressure is required to achieve the stated density, and that as long as the entropy increase during compression keeps the temperature below 5 million degrees, the gas can be considered cold

and the compression process will be efficient. Pressures between 10 and 100 gigabars are representative of conditions required to create fusion in thermonuclear weapons.

A useful rule-of-thumb about electron density in various materials can be obtained by observing that most isotopes of most elements have a roughly 1:1 neutron/proton ratio in the nucleus. Since the number of electrons is equal to the number of protons, we can assume that most substances contain a fixed number of electrons per unit mass: 0.5 moles/gram (3.01 x 10^23 electrons). This assumption allows us to relate mass density to the Fermi gas pressure without worrying about chemical or isotopic composition.

The stable isotopes of most light elements follow this rule very closely, for two that are commonly used as fuel in thermonuclear weapons (Li-6 and D) it is exact. Very heavy elements contain somewhat fewer electrons per gram, by 25-30%. The largest deviations are the lightest and heaviest isotopes of hydrogen: 1 mole/gram for ordinary hydrogen, and 0.333 moles/gram for tritium.

Since the only way a cold Fermi gas can acquire additional energy is in the form of electron kinetic energy, when the thermal energy is substantially above T_Fermi, then the kinetic energy dominates the system and the electrons behave like a classical Boltzmann gas.

Thus as the electronic shells of atoms break down, the value of gamma approaches a limiting value of 5/3 with respect to the total internal energy, regardless of whether it is thermal or quantum mechanical in nature.

The total pressure present is the sum of the Fermi pressure, the kinetic pressure of the Boltzmann gas consisting of the nuclei and non-degenerate electrons, and the pressure of the photon Bose gas. Similarly, the energy density is the sum of the contributions from the Fermi, Boltzmann, and Bose gases that are present.

Now when electrons are stripped from atoms through thermal ionization, we also have an electron gas which is technically a Fermi gas. We rarely consider thermally ionized plasmas to be Fermi gases though, because usually the electron densities are so low that the thermal energy is much greater than the Fermi energy.

An important consequence of this is the phenomenon of "ionization compression". At STP most condensed substances have roughly the same atom density, on the order of 0.1 moles/cm^3; the densities can vary considerably of course due to differing atomic masses. By the rule of thumb above, we can infer that electron densities also roughly mirror mass densities.

If two adjoining regions of STP condensed matter of different electron density are suddenly heated to the same extremely high temperature (high enough to fully ionize them) what will happen?

Since the temperature is the same, the radiation pressure in both regions will be the same also. The contribution of the particle pressure to the total pressure will be proportional to the particle density however. Initially, in the un-ionized state, the particle densities were about the same.

Once the atoms become ionized, the particle densities can change dramatically with far more electrons becoming available for dense high-Z materials, compared to low density, low-Z materials. Even if the system is radiation dominated, with the radiation pressure far exceeding the particle pressures, the total pressures in the regions will not balance. The pressure differential will cause the high-Z material to expand, compressing the low-Z material.

The process of ionization compression can be very important in certain thermonuclear systems, where high-Z materials (like uranium) are often in direct contact with low-Z materials (like lithium hydrides).

It is interesting to note that when matter is in a metallic state, the outermost electrons are so loosely bound that they become free. These electrons form a room-temperature plasma in the metal, which is a true Fermi gas. This electron plasma accounts for the conductance and reflectivity of metals.

3.2.4.1 Thomas-Fermi Theory

A widely used approximate theory of the high pressure equation of state was developed in 1927-1928 that ignores the electron shell structure of matter entirely. Called the Thomas-Fermi (TF) theory, it models matter as a Fermi gas of electrons with a Boltzmann gas of nuclei evenly distributed in it, using a statistical description of how the electron gas behaves in the electrostatic field.

The Thomas-Fermi theory includes only the repulsive forces of the electron gas, and the thermal pressure, and ignores the attractive forces that hold solid matter together. It is thus a good approximation of matter only at high enough pressures that repulsive forces dominate. Fortunately experimental EOS data is available at pressures extending into this range (several megabars). Various adjustments to TF theory have been proposed to extend its range of application, such as the Thomas-Fermi-Dirac (TFD) model that includes attractive forces (others exist - Thomas-Fermi-Kalitkin, etc.)

TF theory was employed at Los Alamos together with the existing high-pressure EOS data (at that time only up to hundreds of kilobars) to perform the implosion calculations for the plutonium bomb. Elements with high electron densities (which, from the above rule-of-thumb, is more or less equivalent to elements with high mass densities) are described reasonably well by the Thomas-Fermi model at pressures above about 10 megabars.

3.2.5 Matter At High Temperatures

If the thermal or kinetic energy of the atoms in a substance exceeds the binding and compressive energies, then regardless of pressure it becomes a gas. In highly compressed condensed matter, this occurs at several tens of thousands of degrees C. When the kinetic energy substantially exceeds the combined energies of all other forms of energy present, matter behaves as a perfect gas.

At sufficiently high temperatures, the outer electrons of an atom can become excited to higher energy levels, or completely removed. Atoms with missing electrons are ions, and the process of electron removal is called ionization. Then number of electrons missing from an atom is its ionization state. Excitation and ionization occurs through collisions between atoms, collisions between atoms and free electrons, and through absorption of thermal radiation photons. When all of the atoms have become ionized, then matter is said to be "fully ionized" (in contrast the phrase "completely ionized" usually refers to an atom that has had *all* of its electrons removed).

The energy required to remove an unexcited electron is called the ionization energy. This energy increases with each additional electron removed from an atom due to the increase in ionic charge, and the fact that the electron may belong to a shell closer to the nucleus. The ionization energy for the first ionization state is typically a few electron volts. Hydrogen has one of the highest first ionization energies (13.6 eV), but most elements have first ionization energies of 4 - 10 eV. The energy required to remove the last electron from a plutonium atom (the 94th ionization state) in contrast is 120 KeV. The first and last ionization energies for some elements common in nuclear weapons are:

Table 3.2.5. Representative First and Last Ionization Energies

Element	First Ionization (eV)	Last Ionization (eV)
Hydrogen	13.598	-
Lithium	5.39	3rd: 122.4
Beryllium	9.32	4th: 217.7
Oxygen	13.61	8th: 871.1
Uranium	6	92nd: 115,000
Plutonium	6.06	94th: 120,000

There is a simple law for computing the ionization energy of the last electron (the Zth ionization state for atomic number Z):

Eq. 3.2.5-1

$E_{i,Z} = Z^2 * 13.6\ eV$

For other ionization states, the other electrons bound to the nucleus provide partial screening of the positive charge of the nucleus and make the law more complex.

The energy required to excite an electron is less than the associated ionization energy (E_i). An excited electron is more easily removed from an atom, the energy required being exactly the difference between the unexcited ionization energy and the excitation energy. Under high pressures, excited electrons are subjected to strong Coulomb forces which tend to remove them from the atom. Also frequent atom-atom, atom-electron, and atom-photon interactions will tend to ionize the more weakly bound excited electron. Even if it is not removed, electrons tend to drop back to their ground state after awhile with the emission of a photon. Excitation is thus unimportant in dense, high pressure gases.

The average ionization state of a gas depends on the ionization energy for each ionization state, the temperature of the gas, and the density. At a temperature T, the average particle is kT. If this value is larger than the ionization energy of an electron attached to an atom, then an average collision will remove it (for hydrogen E_i corresponds to T = 158,000 degrees K). We can thus expect the average ionization state to be at least equal to i, where i is the greatest ionization state with ionization energy less than or equal to kT. In fact ionization can be appreciably greater than this, with higher states for denser gases at the same temperature due to more frequent collisions. If the gas density is comparable to the density of condensed matter, then the energy of i is typically in the order of 3kT - 4kT. At the highest temperatures in fission weapons (50 to 100 million degrees K), uranium and plutonium can achieve ionization states of 80 to 85.

Ionization is a statistical process so often a mixture of ion states is present in varying proportions. At the densities and temperatures encountered here though, the effective spread in ionization states is quite narrow, and we can assume that there will effectively be only one ionization state at a given temperature.

Because kT = 120 KeV at 1.4 billion degrees K, complete ionization of these atoms would normally be expected only in the most extreme conditions of thermonuclear weapons (densities of 200-500, and temperatures of 300-350 million degrees), if at all. Note however, that at these extreme densities the atomic shells breakdown and matter exists as a Fermi gas rendering ionization moot.

Every electron dislodged from an atom becomes an independent particle and acquires its own thermal energy. Since the law E = NkT does not distinguish types of particles, at least half of the

thermal energy of a fully ionized gas resides in this electron gas. At degrees of ionization greater much than 1, the thermal energy of the atoms (ions) becomes unimportant.

Since an ordinary (as opposed to a quantum) electron gas can absorb energy only through kinetic motion, it is a perfect gas with gamma equal to 5/3. Although an electron gas is perfect, there are two processes that tend to drive the effective value of gamma down below 5/3 when considering large increases in internal energy. First, it should be apparent that ionization (and excitation) absorbs energy and thus reduces gamma. The second effect is due simply to the increase in N with increasing ionization state. The larger the number of free electrons, each sharing kT internal energy, the larger the sink is for thermal energy. This second effect tends to overwhelm the absorption of ionization energy as far as determining the total internal energy of the gas, but ionization has very important effects on shock waves (discussed below). These effects are especially pronounced in regimes where abrupt increases in ionization energy are encountered (e.g. the transition from an un-ionized gas to a fully ionized gas; and the point after the complete removal of an electron shell, where the first electron of a new shell is being removed).

At high temperatures and ionization states (several millions of degrees and up), where large amounts of energy are required to dislodge additional electrons, both of these ionization effects can often be ignored since the net increase in electron number (and absorption of ionization energy) is small even with large increases in internal energy, and the energy gap between successive electron shells becomes very large.

At very high temperatures, the effect of radiant energy must be taken into account when evaluating the equation of state. Since the energy present as a blackbody spectrum photon gas increases as the fourth power of temperature, while the kinetic energy increases approximately proportionally with temperature (it would be strictly proportional but for the increase in N through ionization), at a sufficiently high temperature the radiation field dominates the internal energy of matter. In this realm the value of gamma is equal to that of a photon gas: 4/3.

3.3 Interaction of Radiation and Matter

Photons interact with matter in three ways - they can be absorbed, emitted, or scattered - although many different physical processes can cause these interactions to occur. For photons with thermal energies comparable to the temperatures encountered in thermonuclear weapons, the only significant interactions with matter are with electrons. The mechanisms by which photons and electrons interact are conveniently divided into three groups: bound-bound, bound-free, and free-free. Bound-bound refers to interactions in which both the initial and final states of the electron involved are bound to an atom. Bound-free describes interactions in which one state is bound to an atom, and the other is a free electron (it doesn't matter whether it is the initial or final state). Free-free interactions are ones in which the electron remains free throughout.

Now each mechanism of photon interaction can operate in a forward or reverse direction. That is, a process that absorbs a photon can also operate in reverse and cause photon emission (making absorption and emission two sides of the same coin); or a scattering process that removes energy from a photon can also add energy. This principle is called *microscopic reversibility*.

Frequently, when solving practical problems, we would rather not consider each individual mechanism of interaction by itself. It is often preferable to have coefficients that describe the overall optical properties of the medium being considered. Thus we have absorption coefficients, which combine all of the individual absorption mechanisms, emission coefficients, and scattering coefficients. Normally we just talk about absorption and scattering coefficients (designated k_a and k_s) because we can combine the absorption and emission processes, which offset each other, into a single coefficient (calling this combined coefficient an "absorption coefficient" is just a matter of convenience). Henceforth, "absorption coefficient" will include both absorption and emission unless otherwise stated. To characterize the optical properties of the medium (also called the "opacity") with a single number we can use the "total absorption coefficient" (also called the attenuation or extinction coefficient), which is the sum of the absorption and the scattering coefficients.

Since, with few exceptions, the cross section of each mechanism of interaction varies with photon energy, the optical coefficients vary with photon energy as well. If we are dealing with a monoenergetic flux of photons (like laser beams) then we need to have absorption and scattering coefficients for that particular photon frequency. If the flux contains photons of different energies then we must compute overall coefficients that average the spectral coefficients over the spectral distribution.

The process of averaging opacity over the photon frequency spectrum is straightforward only if we can assume that all photon emissions are spontaneous (independent of other photons). This is valid if the medium is optically thin, that is, much less than an absorption mean free path in extent. This type of mean free path is called the Planck mean free path. If the medium is optically thick (much larger in extent than the absorption mean free path) then this assumption is not valid. Due to the quantum statistical behavior of photons, the presence of photons increases the likelihood of photon emission to a value above the spontaneous level. This effect ("stimulated emission") is responsible for both the existence of the blackbody spectrum, and of the phenomenon of lasing. When this effect is taken into account during the averaging process, the result is known as the Rosseland mean free path.

In addition to the total absorption coefficient k, the opacity of a medium can be measured using the mean opacity coefficient K. K is a density normalized measure, given in units of cm^2/g (that is total absorption cross section per gram), while k is the actual opacity of the medium at its prevailing density and is given in units of cm^-1. Although K is normalized for mass density, it is often more informative for our purposes to express it in terms of particle or atom density (cm^2/mole). The coefficients K and k are related to the total photon mean free path (l_phot) by:

Eq. 3.3-1

l_phot = 1/(K*rho) = 1/k

3.3.1 Thermal Equilibrium

In a system in thermal equilibrium the *principle of detailed balancing* also applies. This principle holds that each process is exactly balanced by its exact opposite so that the net state of the system is unchanged.

We have already seen that in an equilibrium system, the intensity of the radiation field and its spectrum are governed by the blackbody radiation laws which hold without reference to the actual mechanisms of photon emission and absorption. This indicates that in a fundamental sense, these mechanisms are basically irrelevant to determining the state of the radiation field in the system. The rates at which they occur are governed by quantum principles so that they always generate a blackbody field. If the optical coefficients as a function of photon energy are known then computing the overall coefficients across the whole radiation spectrum is straightforward.

Now having said this, a little qualification is in order. The mechanisms of absorption and emission can produce local features in the blackbody spectrum. For example, a strong absorption line can create a narrow gap at a particular frequency. The energy missing in this gap will be exactly balanced by the increased intensity in the remainder of the spectrum, which will retain the same relative frequency-dependent intensities of the ideal black body spectrum.

A second caveat is that the blackbody spectrum only applies to systems in thermal equilibrium. Specific mechanisms can dominate non-equilibrium situations, and can occur without significant counterbalance by the reverse process. Laser emission and fluorescent emission are common examples of non-equilibrium processes.

Although the specific interaction mechanisms in an equilibrium system do not affect the radiation spectrum, they still affect the optical coefficients because the photon-matter interaction cross sections, and thus the spectral coefficients, do depend on the mechanisms involved. The physics of these mechanisms is often extremely complex (especially the bound-bound and bound-free processes), even the process of developing simplifying approximations is hard. It is thus often very difficult to determine what the values of the optical coefficients should be.

If a system approximates thermal equilibrium then the blackbody spectrum can be used to compute overall optical coefficients for the medium. The coefficients vary with temperature, not only because the peak of the blackbody spectrum varies with temperature, but because the interaction mechanisms are usually temperature dependent also (and may be density dependent as well).

Approximate thermal equilibrium occurs where the change in the radiation field is gradual with respect to both distance and time (these are necessary and sufficient). Gradual change with distance means that the change in energy density is relatively small over the distance l_phot. Gradual change with time means that the energy density does not change much over a radiation relaxation period. Radiation relaxation is usually so fast that this last condition is of little importance. Since typically the spectral mean free path changes with frequency, it is possible for only a portion of the spectrum to be in equilibrium. If the portions of the spectrum that are not in equilibrium make only a small contribution to the energy density, they can be ignored.

The conditions for thermal equilibrium exist in the interior (at least one optical thickness from the surface) of an optically thick body where the transport process is dominated by scattering. This situation also ensures the existence of local thermodynamic equilibrium (LTE), that is the radiation field and matter are in thermodynamic equilibrium at every point in the medium. LTE guarantees that the spectral distribution of the flux will be a blackbody spectrum, and further that the spectrum at any point will be determined by the temperature at that point.

3.3.2 Photon Interaction Mechanisms

It is useful to survey the mechanisms by which matter and energy interact in more detail, to gain a deeper understanding of the factors that affect opacity. In some systems of interest (such as fully ionized matter) the coefficient can be calculated directly from the underlying interaction mechanisms, in others it may be extremely difficult and require experimental validation.

If fact often it is not practical to calculate opacity values directly from basic principles. In these cases they must be determined by direct measurements in the systems of interest. It is also possible to estimate upper and lower bounding values, but these are often very loose.

3.3.2.1 Bound-Bound Interactions

These interactions occur when an electron bound to an atom (or ion) moves between energy levels as a result of photon capture or emission. Capture raises the electron to a higher energy level, emission causes the electron to drop to a lower one. The phenomena known as fluorescence is the process of capturing a photon, followed by the emission of a lower energy photon as the excited electron drops back to its ground state through some intermediate energy absorbing process.

Photon capture by this mechanism requires that that the photon energy correspond exactly to the energy difference between the energy level an electron is currently occupying, and a higher one. The capture cross section for a photon meeting this criteria is extremely large, otherwise it is zero. Quantum uncertainty gives the absorption line a narrow finite width, rather than the zero width an infinite precision match would indicate.

In principle any atom has an infinite number of possible energy levels. In practice, at some point the binding energy to the atom is so weak that the electron is effectively free. The hotter and

denser matter gets, the lower is the energy level (and the greater is the binding energy) where this occurs.

In hot dense gases that are not completely ionized (stripped of electrons) line absorption contributes significantly to the opacity of the gas, and may even dominate it.

Fermi gases resemble completely ionized gases in that no electron is bound to a nucleus. In both cases bound-bound transitions between atomic quantum energy levels cannot occur. Fermi gases have quantum energy levels of their own however, and bound-bound transition between these energy levels are possible.

3.3.2.2 Bound-Free Interactions

The process in which a bound electron absorbs a photon with energy at least equal to its binding energy, and thereby is dislodged from the atom is called the photoelectric effect. This process can occur with any photon more energetic than the binding energy. The cross section for this process is quite significant. When matter is being heated by a thermal photon flux (that is, the matter is not in thermal equilibrium with the flux), photoelectric absorption tends to dominate the opacity.

The reverse process by which atoms capture electrons and emit photons is called radiative electron capture. The balance between these two processes maintains the ionization state of an equilibrium gas.

3.3.2.3 Free-Free Interactions

There are two principal mechanisms by which photons interact with free electrons. These are a photon emission process called bremsstrahlung (and its reverse absorption process), and photon scattering. Naturally, these are the only processes that occur in a completely ionized gas (atoms are completely stripped). In a highly ionized gas (all atoms are ionized, and most but not all of the electrons have been removed) these processes often dominate the opacity. Theoretically free-free interactions are much easier to describe and analyze than the bound-bound and bound-free processes.

3.3.2.3.1 Bremsstrahlung Absorption and Emission

The term "bremsstrahlung" is German and means "slowing down radiation". It occurs when an electron is slowed down through being scattering by an ion or atom. The momentum, and a small part of the energy, is transferred to the atom; the remaining energy is emitted as a photon. Inverse bremsstrahlung (IB) occurs when a photon encounters an electron within the electric field of an atom or ion. Under this condition it is possible for the electron to absorb the photon, with the atom providing the reaction mass to accommodate the necessary momentum change. In principal bremsstrahlung can occur with both ions and neutral atoms, but since the range of the electric field of an ion is much greater than that of a neutral atom, bremsstrahlung is a correspondingly stronger phenomenon in an ionized gas.

In a bremsstrahlung event, we have:

m_e*d_v_e = m_i*d_v_i

due to the conservation of momentum (m_e and d_v_e are the electron mass and velocity change, m_i and d_v_i are for the ion). Which gives us:

m_i/m_e = d_v_e/d_v_i

Since m_i/m_e (the mass ratio between the ion and electron) is 1836*A (where A is the atomic mass of the ion), the velocity change ratio is 1/1836*A. Kinetic energy is proportional to m*v^2, so the kinetic energy change for the ion is only about 1/1836*A of the energy gained or lost by the electron. Bremsstrahlung/IB is thus basically a mechanism that exchanges energy between photons and electrons. Coupling between photons and ions must be mediated by ion-electron collisions, which requires on the order of 1836*A collisions.

Unlike the bound-bound and bound-free processes, whose macroscopic cross section is proportional to the density of matter, the bremsstrahlung/IB cross sections increase much faster with increasing density. It thus tends to dominate highly ionized, high density matter.

The absorption coefficient k_v from bremsstrahlung (assuming a Maxwellian electron velocity distribution at temperature T) in cm^-1 is:

Eq. 3.3.2.3.1-1

k_v = 3.69 x 10^8 (1 - Exp(-h*nu/k*T))(Z^3 * n_i^2)/(T^0.5 * nu^3)

or

Eq. 3.3.2.3.1-2

k_v = 2.61 x 10^-35 (1 - Exp(-pe/k*T))(Z^3 * n_i^2)/(T^0.5 * pe^3)

where Z is the ionic charge, n_i is the ion density, T is the electron temperature (K), and nu is the photon frequency, h is Planck's constant (6.624 x 10^-27 erg cm^2/sec), k is Boltzmann's constant, pe is photon energy (eV), and the other units are CGS.

We can compute the effective overall absorption coefficient by averaging across the frequency spectrum:

Eq. 3.3.2.3.1-3

$$k_1 = 6.52 \times 10^{-24} (Z^3 * n_i^2)/(T^3.5)$$

The absorption mean free path (in cm)is simply 1/k_1:

Eq. 3.3.2.3.1-4

$$l_1 = 1.53 \times 10^{23} (T^3.5)/(Z^3 * n_i^2)$$

The total amount of energy emitted per cm^3 per sec (assuming electron and photon thermal equilibrium) is:

Eq. 3.3.2.3.1-5

$$e = 1.42 \times 10^{-27} Z^3 n_i^2 T^0.5 \text{ ergs/cm}^3\text{-sec}$$

The mean free path formulae given above is based on the assumption that all photon emissions are spontaneous (i.e. the medium is optically thin). Compensating for stimulated emission gives

us the Rosseland mean free path for bremsstrahlung, which is longer than the spontaneous value by a constant factor:

Eq. 3.3.2.3.1-6

$$l_R = 4.8 \times 10^{24} (T^{3.5})/(Z^3 * n_i^2)$$

3.3.2.3.2 Scattering

All of the processes described so far are absorption and emission processes which involve photon destruction and creation, and necessarily exchange substantial amounts of energy between matter and the radiation field. Photon scattering does not, in general, involve significant exchanges in energy between particles and photons. Photon direction and thus momentum is changed, which implies a momentum and energy change with the scattering particle. But photon momentum is usually so small that the energy exchange is extremely small as well.

3.3.2.3.2.1 Thomson Scattering

The only way a photon can interact with an electron in the absence of a nearby atom or ion is to be scattered by it. In classical physics this scattering process (which is due to the electron acting as a classical oscillator) cannot change the photon energy and has a fixed cross section known as the Thomson cross section, which is 6.65 x 10^-25 cm^2. This is multiplied by the electron density to obtain the scattering coefficient k_s. The scattering mean free path is then:

Eq. 3.3.2.3.2.1-7

$$l_s = 1.50 \times 10^{24} /(Z * n_i)$$

Referring to Eq. 3.3.3.1-3, above we can see that at a sufficiently high temperature, the bremsstrahlung absorption coefficient may become smaller than the Thomson coefficient, which will then tend to control radiation transport.

3.3.2.3.2.2 Compton Scattering

When the photon energy becomes comparable to the electron rest-mass (511 KeV), the photon momentum is no longer negligible and an effect called Compton scattering occurs. This results in a larger energy-dependent scattering cross section. Compton scattering transfers part of the energy of the photon to the electron, the amount transferred depends on the photon energy and the scattering angle. 511 KeV is an energy much higher than the typical thermal photon energies encountered in our domain of interest. At high thermonuclear temperatures (35 KeV) a significant proportion of photons in the upper end of the Planck spectrum will undergo this process.

In a hot gas, where electrons have substantial kinetic energy, inverse Compton scattering also occurs which transfers energy from the electron to the photon.

3.3.3 Opacity Laws

As noted earlier, the actual opacity of a material is the sum of the absorption and scattering coefficients:

Eq. 3.3.3-1

$$k_t = k_a + k_s$$

The total photon mean free path is then:

Eq. 3.3.3-2

$$l_{phot} = 1/k_t = 1/(k_a+k_s) = 1/((1/l_a)+(1/l_s)) = (l_a*l_s)/(l_a+l_s)$$

The absorption and scattering coefficients, k_a and k_s, are in turn the sum of all of the component absorption and scattering processes described in the subsections above. In some cases, not all of these are present to a significant extent, but in others there may be many contributors. In performing photon transport calculations though, it is very undesirable to individually model and compute each interaction process. These processes tend to be rather complicated in form and, at worst, are almost intractably difficult to compute. Simulations that include photon transport as but one component would be exorbitantly expensive computation-

wise if these processes were always included explicitly. And it is completely impossible to conduct analytical studies of transport processes without introducing radical simplifications.

Formulating laws of simpler form is essential to the practical study of photon transport. These laws are approximations that describe the dependency of the total opacity on temperature and density in the range of interest. They are not universally valid and care must be taken to ensure that the law chosen is actually applicable in the range where it is to be applied.

A popular form of an opacity law is:

Eq. 3.3.3-3

$$K = (K_0 * \rho^C * T^{-m}) + K_1$$

where K, K_0, and K_1 are the opacities in cm^2/g, rho is density, T is temperature, and C and m are constants. According to this law, at high temperatures the opacity converges to K_1, while as T approaches zero the opacity goes to infinity (and thus the MFP drops to zero).

Opacity laws are simplest when we can assume complete ionization of the material. For an optically thick medium we can easily derive these factors from Equation 3.3.2.3.1-6 and 3.3.2.3.1-7: m = 3.5;

C= 1; and

K_0 = 2.1x10^-25 * Z^3 * (Avg/A)^2;

K_1 = 6.65x10^-25 * Z * (Avg/A);

where Z is the atomic number, Avg is Avogadro's Number, A is the atomic mass. The first term is due entirely to bremsstrahlung, and the second term (K_1) is due to Thomson scattering.

For fusion fuels, usually mixtures of hydrogen, lithium, and helium, complete ionization can be assumed due to their relatively low ionization temperatures, and this law is applicable. Higher Z materials that may be used in the construction of nuclear weapons, like carbon, aluminum, iron, lead, tungsten, and uranium must take into account the effects of bound electrons. For these higher Z materials an m value of 1.5 to 2.5 is typical. For very high Z materials, like uranium, an m value of 3 is commonly assumed. Since absorption through bound electrons is due to photon-atom (or photon-ion) interactions, for a given ionization state opacity should be roughly proportional to atom (ion) density giving us C = 0. Strict proportionality doesn't occur since the ionization state is affected by density.

The appropriate values of K_0 for high Z materials at the densities and temperatures of interest are not readily available however. In the United States, opacities for elements with Z > 71 are still classified. Even for elements with Z < 71, for which the opacities are not classified, U.S. government data has not yet been released, and little information from other sources appears available in the public domain.

Parameters for the opacity law can be derived for some elements with Z up to 40 from figures published in Los Alamos Report LAMS-1117. This document covers temperatures (0.5 - 8 KeV) and densities (1-30 g/cm^3 for carbon, 3.5-170 g/cm^3 for copper) in the range of interest for radiation implosion systems.

Table 3.3.4-1. Opacity Law Parameters

Element	Z	m	C	K_0	K_1
Carbon	6	4.9	0.92	0.074	0.22
Aluminum	13	5.0	0.93	1.4	0.21
Iron	26	2.9	0.62	6.6	0.15
Copper	29	2.7	0.55	7.0	0.14
Zirconium	40	1.8	0.72	1.2	0.15

Although some general trends seem apparent in the table above, they cannot be extrapolated easily to much higher Z values. We can see that zirconium disrupts trends in C and K_0 indicated by Z=6 through Z=29. As we ascend the periodic table and electron shells are successively filled, and new ones started, abrupt changes in ionization behavior, and hence opacity will occur.

3.3.4 Radiation Transport

The transport of radiation through matter often alters the physical state of matter in significant ways. Matter absorbs and emits radiation which leads to it heating up and cooling down. The existence of significant radiation transport implies the existence of energy density and temperature gradients, which in turn imply the existence of pressure gradients. Given sufficient time, pressure gradients may lead to significant fluid flow, and corresponding changes in the density of matter. Problems that include both the motion of matter and radiation are the domain

of radiation hydrodynamics, which is discussed in another subsection below. This subsection generally treats the properties of matter as being static (not affected by the transport process), although the effects of heating are considered under radiation heat conduction.

When a ray of photons pass through a layer of matter, the initial flux is attenuated according to the law:

Eq. 3.3.4-1

$$I_x = I_0 * Exp[-x/l_phot]$$

The flux attenuates exponentially, decreasing to 1/e over a thickness equal to the photon MFP. The ratio x/l_phot is called the "optical thickness", usually designated "tau". If the optical coefficient is not constant through a layer, then its optical thickness may be found simply be integrating k_t along the photon path.

The energy flux actually deposited in a thin layer of thickness dx located at x_1 is:

Eq. 3.3.4-2

$$I_x = k_a*dx*I_x_1 = k_a*dx*I_0*Exp[-x_1/l_phot]$$

since only the absorption cross section actually deposits energy in the layer.

This model of pure attenuation assumes that photon emission by the layer is not significant. If photon emission does occur then instead of merely attenuating towards zero, the photon flux will converge to an equilibrium value, e.g. the radiation field and matter will relax to state of equilibrium. The photon MFP provides the distance scale required for establishing equilibrium, and accordingly is sometimes called the "relaxation length". Since the velocity of photons is c (the speed of light), the relaxation time is l_phot/c.

An additional limitation of this model which remains even when photon emission is taken into account is that it treats only a "ray" of photons - a beam in which all photons travel in the same direction. In practice, there is always an angular distribution of the photon flux around the net

direction of radiation flow. When the flux drops off rapidly as the angle from the net direction of flow increases, then the "ray" model remains a good approximation.

When there is a substantial flux at large angles from the net direction of flow, then other models of photon transport are needed. Photon fluxes with large angular spreads inevitably arise when the flux travels through a medium that is several optical thicknesses across, since photons are quickly scattered in all directions. Two important approaches to approximating such situations are radiation diffusion and radiation heat conduction.

A final limitation should be noted concerning the "ray" model. For a beam of photons originating outside of the optically thin medium and travelling through it, the spectral distribution of the flux can be anything at all - a monoenergetic flux like a laser beam, a Boltzmann distribution, etc. A single absorption and scattering cross section can only be used if the spectral cross sections are appropriately weighted for the flux spectrum.

3.3.4.1 Radiation Diffusion

Radiation diffusion (based on the application of diffusion theory) describes photon transport where the energy density gradient is small, i.e. in conditions of approximate local thermal equilibrium (LTE). This situation exists in the interior (at least one optical thickness from the surface) of any optically thick body.

The rate of energy is determined by the diffusion equation:

Eq. 3.3.4.1-1

$$S = -((l_{phot} c)/3) * Grad(U)$$

where S is the energy flux (flow rate), c is the speed of light, U is the energy density, and Grad(U) is the energy density gradient (the rate at which energy density changes with distance along the vector J). Due to the gradual change in state in diffusion, if the thickness of the body is x then the flux can be approximated by:

Eq. 3.3.4.1-2

$$S = ((l_{phot} c)/x) * U$$

That is, the flux is inversely proportional to the optical thickness.

During diffusion, photons are scattered repeatedly and advance in a random walk, travelling much farther than the straight-line distance. In the diffusion process a photon undergoes on the order of (x/l_phot)^2 interactions which the material while advancing the distance x. The characteristic diffusion time t_diff is thus (x^2/(c*l_phot)). When the physical scale of a system is much larger than the photon MFP, the photon field is effectively trapped and is always in local equilibrium with the state of the material, and it is possible for the material to flow faster than the photon flux itself. In the case of stars, where l_phot is on the order of a millimeter, but x is on the order of a million kilometers, it can take millions of years for photons created in the interior to diffuse to the surface (they are absorbed and reemitted many times during the journey of course)!

It is interesting to note that the Rosseland mean free path is equal to the spectral mean free path when the energy of photons is given by:

$$h\nu = 5.8\ kT$$

This means that when a body is optically thick, and radiant energy transport occurs through radiation diffusion, the main energy transport is due to photons with energies much higher than the average particle kinetic energy. On the other hand when a body is optically thin, and energy is radiated from the entire volume, then most of the energy is transported by photons with energies only moderately higher than the average kinetic energy.

This phenomenon is almost always true to some extent, and can greatly affect radiation transport. If the MFP increases with energy (as it does with bremsstrahlung in optically thick bodies), then energy transport is due to very energetic photons which may only comprise a small fraction of the radiation field. On the other hand, the converse may be true. This is the case during early fireball growth in nuclear explosion, when only the radiation in the normal optical range can escape to infinity from the X-ray hot fireball.

3.3.4.2 Radiation Heat Conduction

This approximation of radiation transport is very closely related to the diffusion approximation, and like diffusion applies when LTE exists. The principle difference between radiation heat conduction and diffusion is that the conduction equations are based on temperature gradients, rather than energy density gradients. This is important because at LTE, the temperature of matter and radiation are equal, and the physical state and thermodynamic properties of matter are a function of temperature. If we want to take into account the effects of radiation heating on matter then we must treat the problem in terms of temperature.

3.3.4.2.1 Linear Heat Conduction

Before considering radiation heat conduction, it is helpful to first look at ordinary heat conduction - the familiar conduction process encountered in everyday life.

The ordinary conductive heat flux is given by:

Eq. 3.3.4.2.1-1

$$S = -\kappa * \mathrm{Grad}(T)$$

where S in the flux, kappa is the coefficient of thermal conductivity of the material, and Grad(T) is the temperature gradient. Since the conductive flux is proportional to the temperature gradient this is called linear heat conduction.

The basic behavior of linear heat conduction can be observed by considering a uniform medium at zero temperature bounded by a flat surface on one side. If an arbitrarily thin layer at this surface were instantly heated by the application of a pulse of energy E (in erg/cm^2), how would the heat flow through the medium? Initially the temperature at the surface is some finite (but arbitrarily high) temperature T_0. At the boundary of this heated layer the temperature gradient is infinite since it drops from T_0 to zero degrees in zero distance. Since the flux (rate of flow) is proportional to the temperature gradient, the rate of flow is infinite as well. At any point where a non-zero temperature is adjacent to the unheated medium this will be true. Heat thus instantly diffuses throughout the medium, establishing a temperature profile (and temperature gradient) that is highest at the heated surface and drops to zero only at infinity. The rate of flow drops with the temperature gradient though, so at great distances the actual flow is extremely small. The temperature profile assumes a Gaussian distribution in fact, with nearly all the thermal energy located in a region "close" to the surface.

Figure 3.3.4.1-1. Profile of Linear Heat Conduction

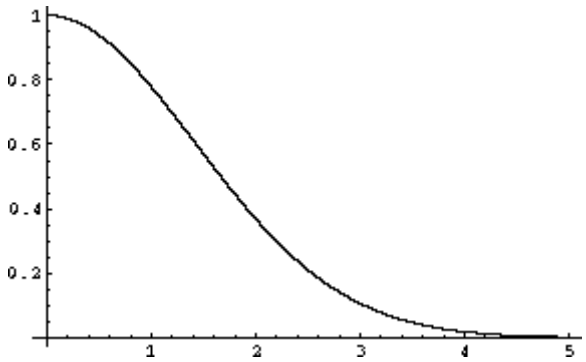

What does "close" mean? The scale of the region is determined by the elapsed time since the temperature jump was introduced, and the thermal diffusivity of the medium (which is usually designated by chi). The thermal diffusivity is given by:

Eq. 3.3.4.2.1-2

$$\chi = \kappa/(\rho * c_p)$$

where kappa is the thermal conductivity of the medium, rho is its density, and c_p is the specific heat at constant pressure. The thermal diffusivity of a gas is approximately given by:

Eq. 3.3.4.2.1-3

$$\chi = (l_{part} * v_{mean})/3$$

where l_part is the particle (molecule, atom, etc.) collision mean free path, and v_mean is the mean thermal speed of the particles.

The actual temperature distribution function is:

Eq. 3.3.4.2.1-4

$$T = Q * Exp[(-x^2)/(4*\chi*t)] / (4\pi * \chi * t)^{0.5}$$

where Q is in deg/cm and is given by:

Eq. 3.3.4.2.1-5

$$Q = E/(\rho c_p)$$

If the process occurs at constant pressure, or:

Eq. 3.3.4.2.1-6

$$Q = E/(\rho c_v)$$

Q provides a scale factor for the heat profile since it is the surface energy density, divided by the ability of an infinitely thin layer to absorb heat.

if constant volume is maintained (c_v is the specific heat at constant volume). From this we can see that the depth of the region where most of energy is concentrated is on the order of:

Eq. 3.3.4.2.1-7

$$x \sim (4\chi t)^{1/2}$$

At a depth of a few times this value the temperature becomes very small. The depth x can thus be taken as a measure of the depth of the penetration of heat.

Of course the prediction of infinite heat propagation velocity (even at infinitesimally small energy flows) is not physically real. In a gas it is actually bounded by the ability of "hot" particles to travel several times the MFP without undergoing collisions, and is thus limited to the velocity of the particles present, and is related to the local velocity of sound (the thermal energies and temperatures involved here are so small that they are utterly insignificant however).

3.3.4.2.2 Non-Linear Heat Conduction

We can easily derive a temperature gradient based equation for radiation transport from the radiation diffusion equation Eq. 3.3.4.1-1 by substituting the blackbody radiation energy density expression (see Eq. 3.1.7-6) for U:

Eq. 3.3.4.2.2-1

$$S = -((l_{phot}*c)/3) * Grad(U)$$

$$= -((l_{phot}*c)/3) * Grad((4*\sigma*T^4)/c)$$

To put this equation in the form of Eq. 3.3.4.2.1-1 we need only compute the appropriate value for kappa:

Eq. 3.3.4.2.2-2

$$kappa = ((l_{phot}*c)/3) * dU/dT = ((l_{phot}*c)/3) * c_{rad}$$

$$= (16 * l_{phot} * \sigma * T^3)/3$$

We can see immediately from the above equation that radiation heat conduction is indeed non-linear: kappa has two components that vary with T. The first component is the photon mean free path, l_phot, which is determined by the (temperature dependent) opacity and drops to zero at zero temperature. The second component, dU/dT, is the radiation specific heat (also designated c_rad):

Eq. 3.3.4.2.2-3

$$c_{rad} = dU/dT = (16 * \sigma * T^3)/c$$

which varies as the third power of T and thus also drops to zero at T=0. One immediate consequence of this is that as T approaches zero, so does the thermal conductivity of the medium. The instantaneous permeation of a medium by heat energy as predicted by the linear conduction model thus does not occur, in non-linear conduction there is always a definite boundary between heated and unheated matter that advances at a finite speed.

The existence of a well defined boundary with a finite speed means that radiation heat conduction creates a "thermal wave" (also called a Marshak wave, or a radiation diffusion wave) with the boundary marking the edge of thermal wave front, and the finite speed being the wave velocity. On the other hand, since radiation conduction is mediated by photons which travel at the speed of light, the actual bounding velocity of thermal waves is much higher than classical kinetic heat conduction which is limited by the speed of sound. The temperature profile of a thermal wave differs markedly from the Gaussian distribution of linear heat conduction. In a thermal wave the temperature in the heated zone is essentially constant everywhere, except near the front where is rapidly drops to zero. The thickness of the zone where this temperature drop occurs is the wave front thickness (see Figure 3.3.4.2.2-1).

Figure 3.3.4.2.2-1. Profile of a Thermal (Marshak) Wave

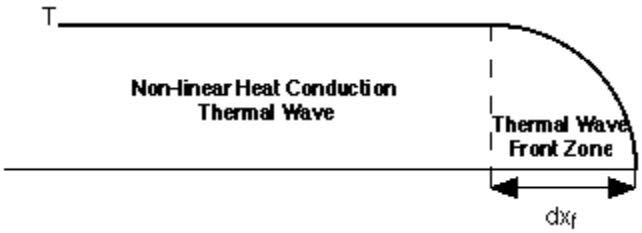

Before we can progress further in studying thermal waves we must obtain expressions for kappa and chi that incorporate the opacity of the medium. We are particularly interested in an expression that is valid at low temperatures, since this is what determines the rate of advance of the heated zone. At low temperatures the opacity becomes very large, so we can neglect the constant factor K_1 in 3.3.-1. The photon MFP is the reciprocal of the opacity-density product, so we can represent l_phot by:

Eq. 3.3.4.2.2-4

$$l_{phot} = 1/((K_0 * \rho^C * T^{-m}) * \rho) = A * T^m, m > 0$$

where A is a constant given by:

Eq. 3.3.4.2.2-5

A = 1/(K_0 * rho^(C+1))

Substituting Eq. 3.3.4.2.2-4 into Eq. 3.3.4.2.2-2 give us:

Eq. 3.3.4.2.2-6

kappa = (16 * A * T^m * sigma * T^3)/3 = (16 * A * sigma * T^n)/3 = B*T^n,

where n = m + 3, and B is a constant given by:

Eq. 3.3.4.2.2-7

B = (16 * A * sigma)/3 = (16*sigma) / 3*(K_0 * rho^(C+1))

In a fully ionized gas where bremsstrahlung dominates n = 6.5, in the multiply ionized region it is usually in the range 4.5 < n < 5.5.

For radiation conduction the expression for chi (Eq. 3.3.4.2.1-2) must use c_v rather than c_p:

Eq. 3.3.4.2.2-8

chi = kappa/(rho * c_v) = ((l_phot*c)/3) * c_rad/(rho*c_v)

= (B * T^n)/(rho * c_v) = a * T^n,

where a is a constant:

Eq. 3.3.4.2.2-9

$$a = B/(\rho * c_v)$$

Note that chi now includes both the specific heat of matter and the specific heat of radiation since two fluids are involved in the transport process.

We can now give an expression for the temperature profile in a thermal wave:

Eq. 3.3.4.2.2-10

$$T = ((n*v/a) * dx)^{(1/n)}$$

This equation gives T where dx is the distance behind the wave front, where n and a are as defined above, and v is the front velocity. This equation indicates that the slope dT/dx is extremely steep when dx is near zero, then flattens out and eventually increases extremely slowly. The scale of the front thickness where this transition occurs is related to the value of (n*v/a). If there some bounding temperature for T, T_1, then we can express the effective front thickness (dx_f) by:

Eq. 3.3.4.2.2-11

$$dx_f = (a * T_1^n)/(n*v) = chi[T_1]/(n*v)$$

Determining exact values for T_1 and v (the position of the front, x_f, as a function of time) require a more complete statement of the problem conditions, but a qualitative expression of the relationship between x_f, time, and the scale parameters Q and a, can be given:

Eq. 3.3.4.2.2-12

$$x_f \sim= (a * Q^n * t)^(1/(n+2)) = (a * Q^n)^(1/(n+2)) * t^(1/(n+2))$$

For the front velocity v:

Eq. 3.3.4.2.2-13

$$v = dx_f/dt \sim= x_f/t \sim= [(a*Q^n)^(1/(n+2))]/[t^((1/(n+2))-1)]$$

$$\sim= (a*Q^n)/X_f^(n+1)$$

These relations show that the front slows down very rapidly. With the parameters a and Q being fixed, the penetration of the front is the (n+2)th root of the elapsed time. If n~=5 then x_f ~ t^(1/7), and v ~ t^(-6/7) ~ x_f^(1/6). That is, the distance travelled is proportional to the seventh root of the elapsed time, and the velocity is the inverse sixth root of distance. Since the total quantity of heat (determined by Q) is fixed, the temperature behind the wave front is a function of how far the front has penetrated into the medium. The driving temperature drops rapidly as the wave penetrates deeper, leading to a sharp reduction in thermal diffusivity (chi), and thus in penetration speed.

It should be noted that if Eq. 3.3.4.2.2-12 is reformulated so that the (temperature dependent) parameter chi replaces Q we get (to within a constant multiple) the same law for the penetration of heat that was seen in linear heat conduction (Eq. 3.3.4.2.1-7):

Eq. 3.3.4.2.2-14

$$x \sim (chi*t)^{(1/2)}$$

A specific situation of interest where Q is fixed is the instantaneous plane source problem described for linear heat conduction - an instantaneous pulse of thermal energy applied to a plane bounding the medium of conduction. In this case an exact solution to this problem is possible, but the result includes a term that is a very complicated function of n (see Zel'dovich and Raizer, *Physics of Shock Waves and High-Temperature Hydrodynamic Phenomena*, for details). Below is the solution equation with an exponential approximation used to simplify it.

Eq. 3.3.4.2.2-15

$$x_f = Exp[0.292 + 0.1504*n + 0.1*n^2]*(a * Q^n * t)^{(1/(n+2))}$$

The only difference between this equation and Eq. 3.3.4.2-13 is the addition of the function of n.

Another interesting case is the problem of radiation penetration into a plane with constant driving temperature. This situation arises physically when the plane surface is in thermal equilibrium with a large reservoir of heat behind it. Examples where this situation include certain phases of fission bomb core disassembly, and certain phases of radiation implosion in thermonuclear weapons. Since the peak driving temperature is constant, the value of chi remains constant, and we get immediately from Eq. 3.3.4.2.2-14:

Eq. 3.3.4.2.2-16

$$x \sim (chi*t)^{(1/2)} \sim (a * T^n * t)^{(1/2)}$$

Under these conditions the radiation wave still decelerates at a respectable pace, but nearly as fast as in the single pulse example. An estimate of the constants for this equation are available for temperatures around 2 KeV (23 million degrees K) when the material is cold uranium from the declassified report LAMS-364, *Penetration of a Radiation Wave Into Uranium*. This is a set of conditions of considerable interest to atomic bomb disassembly, and of even greater interest

to radiation implosion in thermonuclear weapons. When x is in cm, T is in KeV, and t is in microseconds, the equation is:

Eq. 3.3.4.2.2-17

$$x = 0.107 * T^3 * t^{(1/2)}$$

For example, after 100 nanoseconds at 2 KeV the wave will have penetrated to a depth of 0.27 centimeters. Note that this equation implies that for uranium n=6.

If the driving temperature actually increases with time, instead of simply remaining constant, then the radiation wave's propensity to slow down with time and distance will become even more reduced. If we continue to magnify the rate of temperature increase with time, then at some point the radiation wave velocity will remain constant, or even increase with time. As one example of this situation, we can consider the case where the temperature behind the wave front is driven by a constant flux S into the medium. In this situation the temperature behind the wave constantly increases. The temperature (averaged across the heated zone) is approximately:

Eq. 3.3.4.2.2-18

$$T \sim [(S * x_f)/(a * c_v * rho)]^{(1/(n+1))}$$

$$\sim [S^2/(a * c_v * rho)]^{(1/(n+2))} * t^{(1/(n+2))}$$

and the heated zone thickness is:

Eq. 3.3.4.2.2-19

$$x_f \sim [(S^n * (a * c_v * rho)]^{(1/(n+2))} * t^{((n+1)/(n+2))}$$

For n=5, x_f~t^(6/7), T~t^(1/7)~x_f^(1/6), and v~t^(-1/7). This is a very slowly decelerating radiation wave.

An important question is "Under what conditions can the assumption of a stationary medium of conduction be made?". We explicitly assumed that the process of heating the material did not change its density, nor set in motion through the pressure gradients created. We have already seen that a constant density and temperature, there will be a characteristic pressure produced through particle kinetic energy and radiation pressure. We will see in the following section, Section 3.4 Hydrodynamics, that a given pressure and density will produce a pressure wave, which may be a shock wave, with a constant velocity. If the radiation front is decelerating, then at some point, if the medium is thick enough, and the heat transmission process is uninterrupted, this pressure wave must overtake the radiation front. When this occurs, we move into the realm of radiation hydrodynamics proper, discussed in detail in Section 3.5 Radiation Hydrodynamics.

3.4 Hydrodynamics

Hydrodynamics (also called gas or fluid dynamics) studies the flow of compressible fluids, and the changes that occur in the state of the fluid under different conditions. Radiation gas dynamics adds the energy transfer effects of thermal radiation. This becomes significant only when the gas is hot enough for radiation transport to be appreciable compared to the energy transport of fluid flow. Before dealing with the effects of radiation, it is necessary to review the basics of conventional hydrodynamics.

3.4.1 Acoustic Waves

Any local pressure disturbance in a gas that is not too strong will be transmitted outward at the speed of sound. Such a spreading disturbance is called an acoustic wave. This speed, designated c_s, is a function of gas pressure, density (rho) and gamma; or equivalently by the temperature and R (the universal gas constant per unit mass):

Eq. 3.4.1-1

$$c_s = (\gamma * \text{pressure}/\rho)^{0.5} = (\gamma * R * T)^{0.5}$$

Since we have previously dealt with k (Boltzmann's constant), n (particle density), and rho (mass density) but not R is useful to note that:

Eq. 3.4.1-2

$$R = k*n/rho = k*m_{part}$$

where m_part is the average mass of a particle.

If the local pressure disturbance is cyclical in nature, creating alternating high and low pressures, a series of acoustic waves is produced that evokes the sensation of sound. The pressure of sound waves is usually quite weak compared to ambient air pressure, for example even a 130 dB sound (an intensity that inflicts pain and auditory damage on humans) causes only a 0.56% variation in air pressure.

3.4.2 Rarefaction and Compression Waves

A disturbance that decreases gas pressure is called a rarefaction wave, a disturbance that increases pressure is a compression wave. With either type of acoustic wave the change in state of the gas is essentially adiabatic.

When a gas is adiabatically expanded the speed of sound decreases, when it is adiabatically compressed c_s increases. The speed of sound thus decreases behind the leading edge of a rarefaction wave, and increases behind the leading edge of a compression wave. Because of this rarefaction waves tend to stretch out as the propagate, each portion of the disturbance progressively lagging farther and farther behind the leading edge which propagates at the speed of the undisturbed gas. Lagging portions of compression waves tend to catch up with the leading edge since they propagate at higher speeds through the disturbed (compressed) gas. The pressure profile of the compression wave thus becomes steeper and steeper. When a portion of the wave catches up with the leading edge, the slope at that point becomes infinitely steep, i.e. a sudden pressure jump occurs. A wave that causes an instantaneous pressure increase is called a shock wave.

3.4.3 Hydrodynamic Shock Waves

Compression waves are fundamentally unstable. They naturally tend to steepen in time and eventually (if they propagate long enough) will become infinitely steep: a shock wave. On the other hand, rarefaction waves are stable and a rarefaction shock (a sudden pressure drop) is impossible.

3.4.3.1 Classical Shock Waves

The idealized representation of shock waves used in shock wave theory assumes the pressure jump at the front of the wave is a mathematical discontinuity. When a particle of matter passes through an ideal shock front it undergoes an instantaneous change in state. The pressure it is subjected to increases, it is compressed to higher density, its temperature increases, and it acquires velocity in the direction of shock wave travel.

A classical shock wave is distinguished from non-classical shocks by the assumption that the only process of energy transport or change of state present is the shock compression itself.

A diagram of such a shock wave is given below. The subscript 0 designates the initial unshocked state, and 1 designates the shock compressed state.

Figure 3.4.3.1-1. Classical Shock Wave

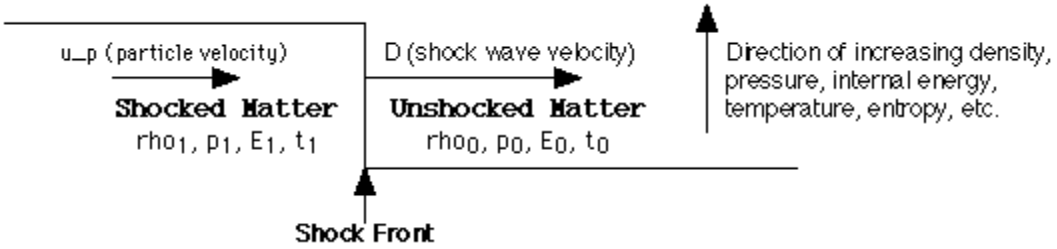

There is a fundamental difference between continuous compression waves and shock waves. The change of state in a compression wave is always adiabatic no matter how steep the pressure profile. It can be shown theoretically that in an instantaneous pressure jump adiabatic compression is impossible, there must be an increase in entropy.

Of course truly instantaneous changes cannot occur in reality. So, just how sudden is this shock transition really? It turns out that the shock transition occurs over a distance which is on the order of the molecular collision mean free path in the substance propagating the shock wave. That is to say, only one or two collisions are sufficient to impart the new shocked state to matter passing through the shock front. In condensed matter, this path length is on the order of an atomic diameter. The shock front is thus on the order of 1 nanometer thick! Since shock waves propagate at a velocity of thousands of meters per second, the transition takes on the order of 10^-13 seconds (0.1 picoseconds)!!

Physically the change in entropy in real shock waves is brought about by viscous forces and heat conduction connected with the discrete nature of matter at the scale of the shock front thickness. Since even very fast continuous compression occurs many orders of magnitude slower than shock compression, with a compression front thickness correspondingly thicker, viscosity is almost always negligible (thus validating the idea of adiabatic compression).

The shock velocity D is always supersonic with respect to the unshocked matter. That is, D is always higher than the speed of sound c_s for the material through which the shock is propagating.

The shock velocity is always *subsonic*with respect to the compressed material behind the front. That is, D - u_p is always less than the speed of sound in the shock compressed material. This means that disturbances behind the shock front can catch up to the shock wave and alter its behavior.

Obviously a shock wave does work (expends energy) as it travels since it compresses, heats, and accelerates the material it passes through. The fact that events behind the shock front can catch up to it allows energy from some source or reservoir behind the shock wave (if any) to replenish the energy expended (a fact of great significance in detonating high explosives).

Mathematically it is often simpler to view a shock front as stationary, and assume that matter is flowing into the shock front at a velocity u_0 equal to -D, and that it is compressed and decelerated to u_1 by passing through the front (u_1 = u_p - D). This is actually the natural way of viewing shock waves in certain situations, like the supersonic flow of air into a nozzle.

It is also often convenient to use the "specific volume" instead of density to simplify the mathematics. Specific volume v_0 is defined as 1/rho_0; v_1 = 1/rho_1.

There is a simple relationship between u_0, u_1, and the density ratio between the shocked and unshocked material (rho_1/rho_0). It can easily be understood by considering the requirement for the conservation of mass. Since matter cannot accumulate in the infinitely thin shock front, the mass flow rate on either side must be the same. That is:

Eq. 3.4.3.1-1

$$u_0 \rho_0 = u_1 \rho_1$$

which gives:

Eq. 3.4.3.1-2

$$u_0/u_1 = \rho_1/\rho_0 = D/(D - u_p)$$

We can gain some intuitive idea of the behavior of a shock wave by examining some limiting cases, very weak and very strong shocks. When we talk about the strength of a shock wave a number of different measures of "strength" can be used. In ordinary gases (i.e. not condensed matter) the pressure ratio p_1/p_0, or the Mach number (D/c_s) are convenient. In either case a very weak shock was a value not much higher than 1, a very strong shock has a value of 10-100 (easy to reach for a gas at atmospheric pressure).

Since the "initial pressure" in condensed matter cannot be measured directly, we use the absolute pressure increase instead of pressure ratios for measuring shock strength. Absolute shock velocity, or more rarely Mach number, is also used. To reach similar limiting "strong shock" conditions for condensed matter, shock velocities of >30 km/sec, and pressures of 10-100 megabars are required, and are beyond the reach of chemical explosives (maximum values are about 400 kilobars and 9 km/sec).

A very weak shock wave travels only very slightly faster than sound, it creates a small pressure increase behind the front, imparts a small velocity to the material it passes through, and causes a extremely small entropy increase. The velocity of sound behind a weak shock is about the same as it is in front. The behavior of a weak shock wave is thus not much different from an acoustic wave (in fact an acoustic wave can be considered to be the limiting case of a weak shock).

As shock strength increases, the proportions by which the energy expended by the wave is divided between compressive work, entropic heating, and imparting motion also changes. The density of matter crossing the shock front cannot increase indefinitely, no matter how strong the shock. It is limited by the value of gamma for the material. The limiting density increase is:

Eq. 3.4.3.1-3

$$\text{density_ratio} = \rho_1/\rho_0 = v_0/v_1 = (\gamma + 1)/(\gamma - 1)$$

Consequently, for a perfect monatomic gas the limit for shock compression is 4. Material with higher gammas (e.g. essentially all condensed matter) has a lower compressive limit. In regimes where abrupt increases in ionization energy absorption occur (initial ionization, initial electron removal from a new shell), the effective gamma can be driven down to the range of 1.20 to 1.25, leading to density increases as high as 11 or 12. As shocks continue to increase in strength above such a regime the kinetic excitation again dominates, and the value of gamma rises back to 5/3 (until the next such regime is encountered).

There is also a limiting ratio between the shock velocity D and u_p:

Eq. 3.4.3.1-4

$$u_p = D*(density_ratio - 1)/density_ratio$$

For a perfect gas this is 3/4 of the shock velocity.

In a strong shock, the limiting partition of energy between kinetic energy and internal energy (compression plus entropy) in the compressed material is 50/50. Since the degree of compression (and thus compressive work) has a fixed limit, entropy must increase without limit. In an extremely strong shock (one where the limiting ratios are approached), almost all of the energy expended goes into kinetic energy and entropic heating.

There are three basic equations, called the Rankine-Hugoniot shock relations (jump conditions), that taken together define the classical shock transition. They are easily derived from conservation laws, in fact one of them has already been given (using the conservation of mass):

Eq. 3.4.3.1-5a Rankine-Hugoniot Relation 1a

$$u_0*rho_0 = u_1*rho_1$$

or

Eq. 3.4.3.1-5b Rankine-Hugoniot Relation 1b

$$D*rho_0 = (D - u_p)*rho_1$$

The mass entering the shock front in a unit of time dt (equal to A*dt*u_0*rho_0) undergoes a velocity change of u_0-u_1, and thus a momentum change of change A*dt*u_0*rho_0*(u_0-

u_1). Conservation of momentum requires that the momentum change be balanced by the force acting over the time dt (momentum is force*time). This force is (p_1 - p_0)*A*dt. We can then derive the second shock relation:

Eq. 3.4.3.1-6 Rankine-Hugoniot Relation 2

$$p_1 - p_0 = u_0*rho_0*(u_0-u_1) = D*rho_0*u_p$$

The kinetic energy entering the shock front in a unit of time dt is equal to mass*(u_0^2)/2, or A*dt*(u_0^3)*rho_0/2. The conservation of energy requires that the energy exiting the front be the same, divided between the kinetic energy leaving the front and the change in internal energy. The kinetic energy leaving the shock front is obviously mass*(u_1)^2/2, and the energy converted to internal energy must then be mass*(u_0^2 - u_1^2)/2. Dividing by the mass we get the third shock relation:

Eq. 3.4.3.1-7 Rankine-Hugoniot Relation 3a

$$E_1 - E_0 = 1/2 (u_0^2 - u_1^2) = 1/2 (2*u_p*D - u_p^2)$$

where (E_1 - E_0) is the change in internal energy per unit mass. Another useful form of this relation, obtained by making the appropriate substitutions, is:

Eq. 3.4.3.1-8 Rankine-Hugoniot Relation 3b

$$E_1 - E_0 = 1/2 (p_1 + p_0)(v_0 - v_1)$$

R-H relations 1 and 2 are independent of the material being shocked. Solving the third relation requires having an equation of state for the material. The form given in Eq. 3.4.3.1-8 is convenient when we are dealing with traditional equations of state, which are expressed in

terms of pressure and volume. For an EOS suitable for use with Eq. 3.4.3.1-7, see Subsection 3.4.3.3 below (Linear Equation of State for Shock Compression).

If we use the gas law EOS, we can write an expression for specific volume as a function of the initial state of the material (p_0, v_0), gamma, and the shock pressure (p_1):

Eq. 3.4.3.1-9

$$v_1 = [p_0*v_0*(gamma + 1) + p_1*v_0*(gamma - 1)]/$$

$$[p_0*(gamma - 1) + p_1*(gamma + 1)]$$

Plotted on the a P-V diagram (pressure vs. volume, see 3.1.5 above) this function defines a Hugoniot curve. This curve is the locus of all possible shock states (produced by shocks of varying strength) from a given initial state. The shock velocity D happens to be equal to the slope of the straight line (called the Rayleigh Line) connecting state 0 and state 1 on the Hugoniot. Relation 1b can then give the particle velocity u_p.

For convenience, D can be calculated from:

Eq. 3.4.3.1-10

$$D^2 = v_0^2 * (p_1 - p_0)/(v_0 - v_1)$$

Or, if we are dealing with a gas we can include the equation of state and omit v_1 explicitly:

Eq. 3.4.3.1-11

$$D^2 = (p_1(1 + gamma) - p_0(1 - gamma))v_0/2$$

The Hugoniot curve lies above the isentropic curve defined by:

Eq. 3.4.3.1-12

p_1*(v_1^gamma) = p_0*(v_0^gamma),

although the two curves are very close together for weak shocks.

Below are curves for isentropic compression, and the shock Hugoniot for a perfect gas with gamma = 5/3. Note that the Hugoniot is very close to the isentrope at low pressures and volume changes, but diverges rapidly as V/V_0 approaches 0.25 (the limiting shock compression).

Figure 3.4.3.1-2. Isentrope and Hugoniot Curves for a Perfect Gas

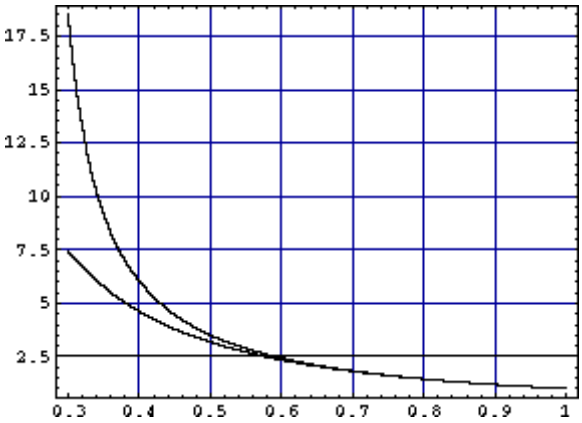

Since force*distance defines work done, and any two points in the P-V diagram define changes in pressure and volume, the area under a curve connecting two points gives the energy change involved in moving between the two states. Expressed in differential form this is:

Eq. 3.4.3.1-13

dE = -p * dv

Adiabatic changes move along the isentropic curve, and the energy change is the area below it. Shock jumps however move instantaneously between the two points, and the energy change must be calculated by the finite difference form:

Eq. 3.4.3.1-14

$$\Delta E = -(p_1 + p_0)/2 * \Delta v.$$

This is the area under the Rayleigh Line connecting the two states. The difference between this area and the area under the isentropic curve passing through state 0 and p_1 gives the increase in entropy due to shock compression.

Here is an example of some actual shock Hugoniot data. It is a plot of the full Los Alamos Scientific Laboratory Shock Hugoniot data set for uranium (from *LASL Shock Hugoniot Data*, ed. S.P. Marsh, 1980) supplemented with a few higher pressure data points found in Skidmore, I.C. and Morris, E., "Experimental Equation of State Data for Uranium and its Interpretation in the Critical Region" (*Thermodynamics of Nuclear Materials*, 1962). The pressure is in kilobars, and goes up to 6.45 megabars. For convenience (mine), the orientation of the plot differs from that traditionally used for Hugoniot curves (such as the perfect gas Hugoniot above). The raw data is plotted with two approximation functions:

Eq. 3.4.3.1-15

For $p < 165$ kilobars

$$V/V_0 = 1 - 0.00059916*p$$

For 165 kilobars $< p <$ 6.45 megabars

$$V/V_0 = 0.68632 + 0.28819*\mathrm{Exp}[-p/600] - 0.000027021*p$$

(with pressure p in kilobars) that fits it quite accurately.

Figure 3.4.3.1-3. LASL Data Set for Uranium

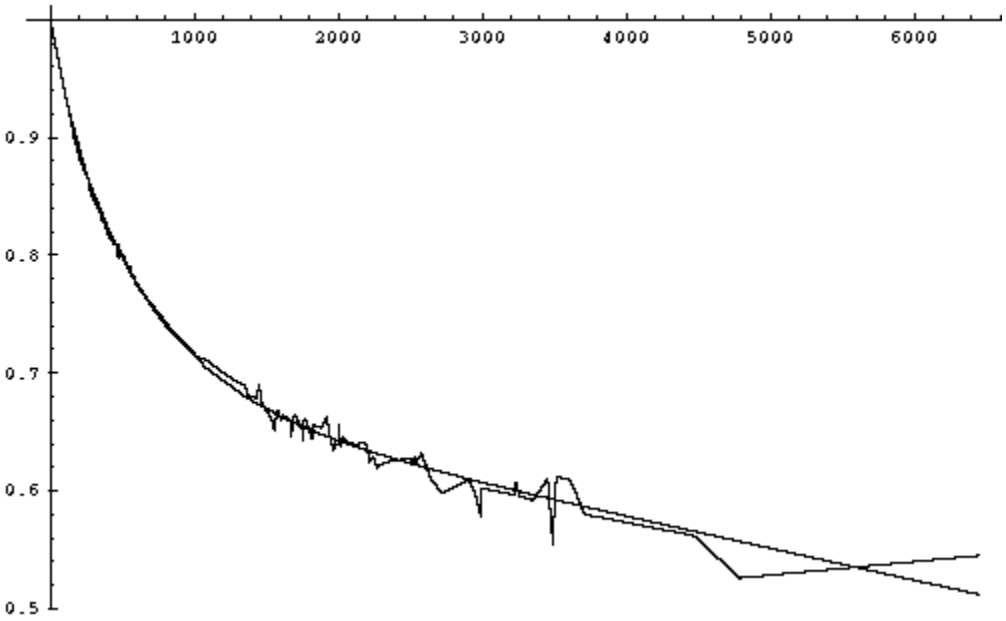

3.4.3.2 Detonation Waves

A detonation wave is a shock wave in an explosive that is supported by the chemical energy released by reactions triggered by the change of state at the shock front. The ability of an explosive to support a detonation wave is the definition of "high explosive" (HE). Low explosives, like black powder ("gun powder") and smokeless powder, merely burn very fast without producing true shock waves.

The chemical reaction in a high explosive takes much longer than the shock transition. The time scale is typically between 1 nanosecond and 1 microsecond, and is thus 1,000 to 1,000,000 times slower than the shock jump. The thickness of the reaction zone thus ranges from on the order of several microns to several millimeters. The ability for the energy released by the high explosive reaction products to reach the detonation front, and thus maintain it, is due to the fact that shock front propagation is subsonic with respect to the shocked explosive.

A given explosive has a characteristic detonation velocity, D. This is the steady state velocity of a detonation wave travelling through the explosive. Representative values for the explosives of major interest in the design of nuclear weapons are given below, density and form are given for each measured velocity:

Table 3.4.3.2-1. Detonation Velocities of Explosives		
Explosive	D (m/sec)	Density/Form
HMX (cyclotetramethylene tetranitramine)	9110	1.89/pressed
RDX/cyclonite (cyclotrimethylenetrinitramine)	8700	1.77/pressed
PETN (pentaerythritol tetranitrate)	8260	1.76/pressed
Composition B (63 RDX/36 TNT/1 wax)	7920	1.72/cast
TATB/TATNB (triamino trinitrobenzene)	7760	1.88/pressed
DATB/DATNB (diamino trinitrobenzene)	7520	1.79/pressed
TNT (trinitrotoluene)	6640	1.56/cast
Baratol (76 barium nitrate/24 TNT)	4870	2.55/cast
Boracitol (60 boric acid/40 TNT)	4860	1.59/cast(?)
Plumbatol (70 lead nitrate/30 TNT)	4850	2.89/cast

The detonation velocity is closely related the energy density of the explosive. Generally the more powerful the explosive, the higher the detonation velocity. Not surprisingly, higher detonation velocities correspond to higher detonation pressures.

Shock waves travelling at velocities higher than the steady detonation velocity (called a strong detonation) will tend to slow down with time to the steady state value (called a Chapman-Jouget, or CJ, detonation). Shock waves that are slower than the detonation velocity (a weak detonation) will tend to accelerate (as long as they are strong enough to initiate the explosion process to begin with).

Besides the energetic reactions of the explosive, expansion of the compressed reaction product gases behind the detonation front also affect the detonation velocity. If the detonation wave is propagating down a cylinder of explosive, gases can expand out to the side immediately behind the wave front. This robs energy from the detonation wave and tends to slow it down. The narrower the cylinder, the greater the reduction in detonation velocity. If the cylinder is narrow enough, the detonation front cannot be sustained and it dies out entirely. In most explosives the drop in velocity becomes significant when the diameter is in the range of 1-5 cm, with failure occurring with diameters in the range of 0.25-1.5 cm. A notable exception to this is PETN, which does not suffer velocity decreases until the diameter is below 0.2 cm, with failure at only 0.08 cm (HMX does nearly as well). For this reason PETN is used to make a flexible detonating cord called "primacord". The velocities in the table above were generally measured using 2.5 cm cylinders, and thus may not represent the maximum CJ detonation velocities.

Some insight into the detonation process can be gained by considering a detonation beginning at the rear surface of an unconfined block of explosive which is wide enough for the edge effect to be negligible. Also assume the detonation front is a plane wave (not curved). The explosive is accelerated in the forward direction by passing through the detonation front, in other words, it gains forward momentum. The conservation of momentum requires that there be a balancing rearward flow. This is brought about by the expansion of the detonation gases out the rear surface of the explosive (a process known as escape).

The expansion of the reaction products (and a corresponding drop in pressure) actually begins as soon as the explosive passes through the shock front. As expansion proceeds with increasing distance from the shock front, the pressure steadily drops, and the internal energy of the gases is progressively converted to kinetic energy directed opposite to the shock wave direction of travel. The rear surface of the reaction zone (called the CJ surface) is defined by the point where the expansion has accelerated the gases rearward enough that they are travelling at exactly the local speed of sound relative to the shock front (energy released behind this point cannot reach the front). Behind this surface the gases (which are still moving forward) continue to expand and decelerate until a point is reached where they are motionless, from this point on the gases move rearward with increasing velocity.

The region of decreasing pressure behind the shock front is sometimes called the Taylor wave.

In a plane detonation propagating through a solid high explosive, the thickness of the zone of forward moving gases grows linearly with detonation wave travel, and is about half the distance the detonation wave has traversed since initiation.

Figure 3.4.3.2-1. Detonation Wave Profile

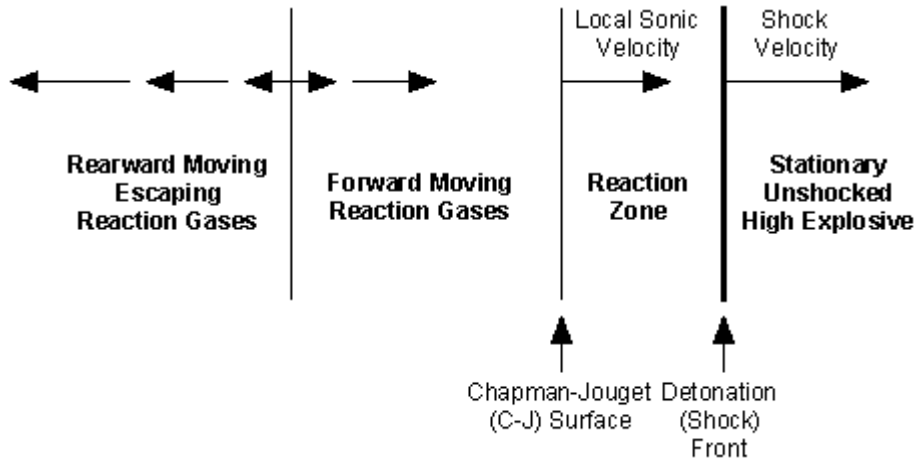

3.4.3.3 Linear Equation of State for Shock Compression

There is a specialized but very useful equation of state that describes the behavior of materials under shock compression. This is called the Linear EOS or the Linear Hugoniot. This EOS describes a very simple relationship between the velocity of the shock wave, and the particle velocity behind the shock:

Eq. 3.4.3.3-1

$$D = C + Su_p$$

where D is the shock velocity, u_p is the particle velocity, C is the speed of sound in the material at zero pressure, and S is a constant for the particular material.

Conservation of mass requires that the degree of compression, rho/rho_0, be equal to:

Eq. 3.4.3.3-2

$$\rho/\rho_0 = D/(D - u_p)$$

Substituting the Linear EOS for U_p, with appropriate rearrangements gives us:

Eq. 3.4.3.3-3

$$\rho/\rho_0 = S/(S + C/D - 1)$$

This implies that the limiting shock compression as D becomes large compared to C is:

Eq. 3.4.3.3-4

$$\rho/\rho_0 = S/(S - 1)$$

The Linear EOS turns out to be quite accurate for virtually all materials under a wide range of conditions, reaching into extreme shock pressures. That this is so is really quite surprising. A wide variety of physical processes occur as materials are shocked from STP to extreme pressures and temperatures. In particular, it will be recalled from section 3.2.2 that the compressibility of the material (indicated by gamma) steadily declines with pressure. No strong theoretical basis exists to explain the linear relationship, or the value of S.

A recent LANL report on this (*General Features of Hugoniots*, by J.D. Johnson, LA-13137-MS, April 1996) shows that the linear range extends up to the astonishing pressure of ~5 gigabars, and shock velocities of ~200 km/sec.

A correspondence can be drawn though between the limiting density given by the linear EOS, and the limiting density of shock compression in a perfect gas:

Eq. 3.4.3.3-5

$$\rho/\rho_0 = S/(S - 1) = (\gamma + 1)/(\gamma - 1)$$

Which implies:

Eq. 3.4.3.3-6

gamma = 2S - 1, or S = (gamma + 1)/2

If we use the value of gamma for a perfect monatomic gas, 5/3, we get:

Eq. 3.4.3.3-7

S = 4/3 = 1.333

Interestingly, the value of S for the elements varies from about 0.9 to 1.9 (most between 1.1 and 1.6), with the distribution centered around 1.33.

Deviations in the linear relationship occur when significant phase changes occur. These show up as kinks in the slope, but affect the overall trend only slightly.

Table 3.4.3.3-1. Linear Equation of State Parameters

Material	C (m/sec)	S	rho_0
Li-6 Deuteride	6590	1.09	0.789
Beryllium	7990	1.124	1.845
Lithium	?	1.15	0.534
Uranium Hydride	2190	1.21	10.92

Tungsten	4040	1.23	19.24
Thorium	2180	1.24	11.68
Perfect Monatomic Gas	-	1.333	-
Uranium	2510	1.51	18.93
Lead	2050	1.52	11.34
Gold	3260	1.56	19.24
Mercury	1800	1.75	13.54
Iron	3770	1.9	7.84
Plutonium	2300	?	19.80

3.5 Radiation Hydrodynamics

Many elements of radiation hydrodynamics of have already been discussed: the properties of blackbody radiation and photon gases, and the mechanisms of interaction between radiation and matter. This section focuses specifically on radiating fluid flow, in particular the limiting case of compressive flows - shock waves, and on how radiation can give rise to hydrodynamic flow.

3.5.1 Radiative Shock Waves

Classical shock waves assume the only process that affects the state of the shocked material is the compression jump at the shock front. In an extremely strong shock the temperature of the material behind the front becomes high enough to radiate significantly, and the thermal radiation transport of energy must be taken into account. The key distinguishing feature of all radiative shock waves is that the radiation of the shocked gas preheats the matter ahead of the kinetic shock jump. Conditions in radiative shock waves are always extreme enough to allow us to treat the shocked material as a gas, and usually a perfect gas.

Classical shocks can be quite luminous of course, but the actual energy flux from radiation is insignificant compared to the hydrodynamic energy flux, resulting in negligible preheating. In air and most other gases the preheating from a classical shock wave is virtually zero because the gas is also transparent to the emitted radiation.

High explosives are incapable of generating radiative shocks. Radiative shock waves arise in collisions at celestial velocities (meteors, comets, etc.), astrophysics (stellar dynamics), and nuclear explosions.

Radiative shock waves can be conveniently divided into four classes in order of increasing strength:

1. Subcritical Shocks
2. Critical Shocks
3. Supercritical Shocks
4. Radiation Dominated Shocks

Ionization effects occur in radiative shock regimes. Initial ionization occurs in the subcritical region. Additional shell ionization effects occur in the supercritical region, and (depending on the element involved) perhaps in the subcritical region as well.

3.5.2 Subcritical Shocks

In these shocks the preheating is relatively small. More precisely, the temperature of the preheated gas is always less than the final (equilibrium) temperature of the gas behind the shock front. An illustrative shock temperature profile is given below.

Figure 3.5.2-1. Subcritical Shock Wave

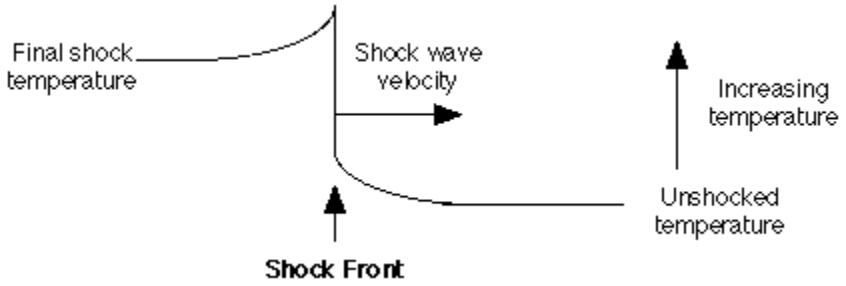

We can see that a small amount of preheating occurs immediately ahead of the shock. The thickness of the preheating layer is approximately equal to the average photon mean free path in the cold gas. The shock jump heats the gas further, creating a temperature spike just behind the front. This layer of superheated gas cools by radiation to the final shock temperature, creating the preheated layer in the process. The superheated gas layer is also about one photon mean free path in thickness.

In air at STP, radiation transport begins to be significant when the shocked gas temperature reaches some 50,000 degrees K. This corresponds to a shock velocity of 23 km/sec (Mach 70), and a preheating temp of 6300 K. The preheating layer will be only 1 mm or so thick.

Subcritical radiative shocks are strong enough to achieve compressions close to the limiting value of (gamma + 1)/(gamma - 1); a factor of 4 for perfect monatomic gases. In strong subcritical shocks in low-Z gases (e.g. air) an effective gamma value of 1.25 is typical due to energy absorption in ionization, resulting in a density increase of 9.

The temperature increase in the preheating layer increases the gas pressure above its initial state, and the resulting pressure gradient generates a weak compression wave in front of the shock. The density increase ahead of the shock front is always small in subcritical shocks. For example, a Mach 70 shock in air results in a density increase on the order of 1% ahead of the front. Most of the compression occurs in the shock jump, but it continues to increase as the superheated gas cools to the final temperature (accompanied by a proportionality much smaller pressure increase as cooling proceeds).

The relationship between the temperature and density at different points in a radiative shock wave, and the final temperature and density, is given by the general law:

Eq. 3.5.2-1

$$T/T_{final} = eta(1- eta)/(eta_{final} * (1 - eta_{final}))$$

where T_final is the final temperature; eta is the volume ratio V_shock/V_0 (the reciprocal of the density ratio); and eta_final is final (and maximum) value of shock compression for the gas in question. This shows that when the preheating is small, so is the compression. When the shock jump compresses the gas to a density close to the final value, the peak temperature is not much higher than the final one.

These effects increase as shock strength comes closer to the critical value (a shock temp of 285,000 K for air at STP). When the final shock temp is 100,000 K (D = 41 km/sec), preheating reaches 25,000 K; at 150,000 K (D = 56.5 km/sec), preheating is 60,000 K; at 250,000 K (D = 81.6 km/sec), preheating is 175,000 K. Above about 90,000 K the preheating layer begins thick enough that it obscures the superheated shock in visible light, an observer in front of the shock sees only the lower temperature preheated layer.

hen the radiation energy density can be neglected, T_final can be calculated from the equation:

Eq. 3.5.2-2

T_final = (2*(gamma - 1)*rho*D^2)/(n*k*(gamma + 1)^2)

where n is the particle density.

Applying the above equation to regimes where strong ionization effects occur presents problems. Ionization increases the particle density, and absorbs additional energy thus driving down the effective value of gamma. If pre-shock values of gamma and n are used in these cases the calculated value of T_final may be greatly overestimated. In ionization regimes closed form equations are not satisfactory and numerical methods are required for exact answers. It also fails close to critical shock conditions when the radiation energy intensity can no longer be neglected. This failure occurs around the point where T_final reaches 70% of T_critical.

Overall, subcritical shocks are not notably different from classical shocks. The energy transport is still dominated by the shock compression, and the equations governing the shock behavior deviate little from the classical equations.

3.5.3 Critical Shock Waves

At critical strength, the radiative preheating equals the final shock temperature. At this point the radiation energy flux equals that of the hydrodynamic energy flux. Above this point we enter the regime of radiation hydrodynamics proper. The increasing dominance of radiation transport now begins to approximate the conditions of thermal waves (radiation diffusion or Marshak waves) discussed in Section 3.3.4.2.2 Non-Linear Heat Conduction. Here the driving temperature for the Marshak wave is provided by the superheated layer behind the shock front, and the Marshak wave itself is the preheating layer in front of the shock.

Figure 3.5.3-1. Critical Shock Wave

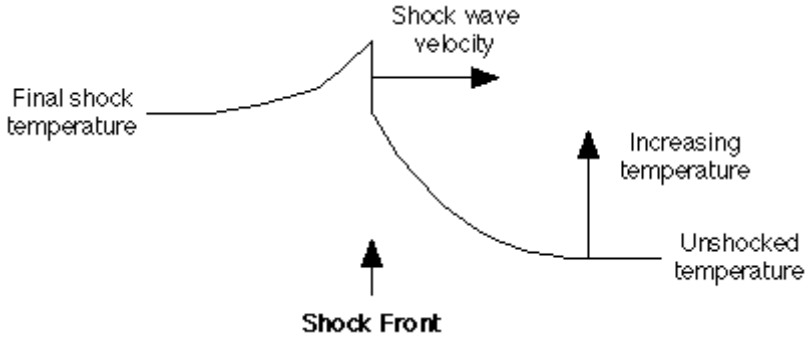

In air initially at STP, critical strength and is reached at a temperature of 285,000 K (T_critical). This corresponds to a shock velocity of 88.1 km/sec. It is this temperature and shock condition

that marks the moment of "hydrodynamic separation" encountered during the cooling of the early fireball of a nuclear explosion, where the hydrodynamic shock takes over from radiation flux in driving the fireball expansion, and causes the brightness of the fireball to fall from its first peak to a temporary minimum.

Determining the exact value of the critical shock velocity and temperature generally requires numerical techniques.

The thickness of the preheated layer becomes greater than a mean free path as shocks near critical strength since the energy flux from the preheated layer becomes so large that it contributes to preheating air even further ahead. The density increase ahead of the shock front is still small.

3.5.4 Supercritical Shock Waves

In these shocks the principal transport of energy is carried out by radiation through the leading Marshak wave. Nearly all of this energy transport is completed before the arrival of the hydrodynamic shock front. The compression wave in the preheated region is still relatively weak however, virtually all of the compression occurs as matter crosses the shock front. Recall though, that in shocks of this strength the energy expended in compression is negligible compared to that expended in heating and shock acceleration.

Figure 3.5.4-1. Supercritical Shock Wave

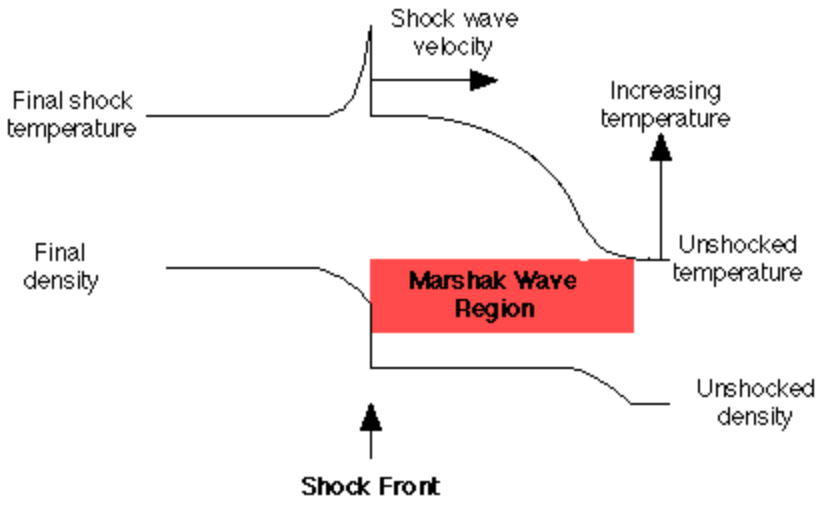

The preheating temperature is the same as the final temperature in all supercritical shocks. As shock intensity increases above critical strength the preheated layer becomes progressively thicker, while the superheated layer becomes thinner. So in a supercritical shock the preheating layer is typically several free paths thick, while the superheated shock layer is less than one free path thick. This is because of the T^4 relationship with radiation flux and radiation density, which causes the shock superheated layer to cool extremely rapidly. The maximum temperature behind the shock wave is:

Eq. 3.5.4-1

$$T_{shock} = T_{final} * (3 - gamma)$$

or 4/3 final for a perfect gas. Since the preheating zone is optically thick, energy transport ahead of the shock front is governed by radiation heat conduction.

Since the shock front temperature is constant, we have here the situation of radiation transport with constant driving temperature governed by equations Eq. 3.3.4.2.2-16 and 3.3.4.2.2-17 (3.3.4.2.2 Non-Linear Heat Conduction). According to these laws, the Marshak wave decelerates rapidly as it moves ahead of the driving temperature front. Because of this there must be a point where the Marshak wave velocity has slowed to exactly the same velocity as the shock front, this point determines the thickness of the preheating layer.

In air, a shock with a final temperature of 750,000 K has a preheated layer some 14 mean free paths thick (about 2 m). At 500,000 K the layer is only 3.4 MFPs thick (40 cm).

It is a supercritical shock that is responsible for the first flash of a nuclear explosion. Even though the fireball is only a few meters across in this phase (for a kiloton range explosion), and even though air is transparent only to an extremely small fraction of the emitted photons, this is still sufficient to make the early fireball as bright (or brighter) at a distance as the later fireball at its maximum size.

3.5.5 Radiation Dominated Shock Waves

In extremely large shocks the radiation pressure and energy density exceeds the kinetic pressure and energy of the gas. At this point we basically have a shock in a photon gas. We might say that the gas particles (ions and electrons) are just "along for the ride". Accordingly, it is the properties of a photon gas (gamma = 4/3) that dominates the situation. The maximum shock compression is thus: (4/3 + 1)/(4/3 - 1) = 7.

In uranium at normal density (18.9), as might be encountered in the core or tamper of an implosion bomb near maximum power output, the point where radiation and kinetic energy densities are the same is reached at 4.8 x 10^7 degrees K. In a lithium deuteride Fermi gas at a density of 200 (typical of compressed fusion fuel) this point is 1.35 x 10^8 K.

In radiation dominated shocks the preheating effect is so large that the one hallmark of classical shock waves - the sudden jump in pressure and density at the hydrodynamic shock front - diminishes and completely disappears in a sufficiently strong shock. This point is finally reached when the ratio between radiation pressure and gas pressure is 4.45. This corresponds to a

temperature of 1.75 x 10^8 degrees K in U-238 (density 18.9), or 2.2 x 10^8 K in LiD (density 200).

Since fission bomb core temperatures can exceed 100 million degrees, and fusion fuel can exceed 300 million degrees, we can expect radiation dominated shock waves to be the normal means of energy transport through matter within a bomb. The temperatures can be expected to cool to being simply a supercritical shock about the time it escapes the bomb casing.

The absence of a sudden pressure jump indicates that the hydrodynamic component of a radiative shock wave has completely disappeared. In this case the radiation dominated shock wave is a pure Marshak wave setting in motion hydrodynamic flow.

Figure 3.5.5-1. Radiation Dominated Shock Wave (Marshak Wave)

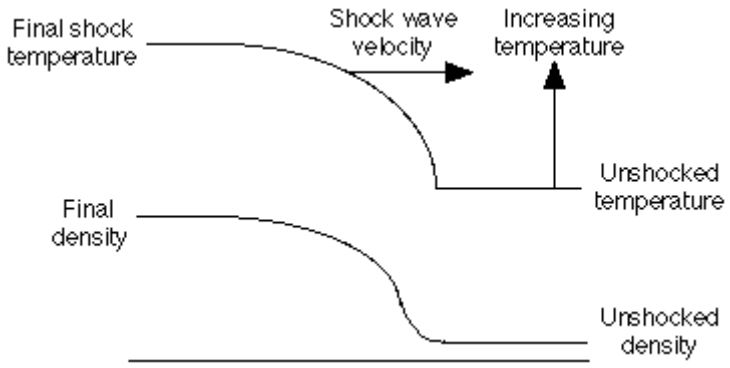

The shock velocity relationship Eq. 3.4.3.1-9 given for classical shock waves: D^2 = v_0^2 * (p_1 - p_0)/(v_0 - v_1) is still valid. By replacing specific volumes with initial density and eta_max; and substituting the sum of kinetic pressure and the blackbody radiation pressure for (p_1 - p_0) we get:

Eq. 3.5.5-1

$$D_{rad}^2 = ((4\sigma/3c)T^4 + nkT)/(\rho(1 - \eta_{max}))$$

$$= ((\gamma + 1)((4\sigma/3c)T^4 + nkT))/(2\rho)$$

where k is the Boltzmann's constant (1.380 x 10^-16 erg/degree K), sigma is the Stefan-Boltzmann constant (5.669 x 10^-5 erg/sec-cm^2-K), and T is temperature (degrees K), c is the speed of light, and n is particle density. Since this equation includes the combined energy

densities of thermal radiation and particle kinetics, it can be used for other radiative shocks as well.

The speed of sound in radiation dominated matter is:

Eq. 3.5.5-2

$$c_{s_rad} = [\gamma((4\sigma/3c)T^4 + nkT)/\rho]^{0.5}$$

Using this shock velocity relationship we can calculate the shock velocity in a material such as uranium. At 10^8 degrees K it is 1480 km/sec; at 4.8×10^7 K it has slowed to 533 km/sec; at 10^7 K it is a mere 166 km/sec. Although 1480 km/sec is very fast, being 160 times higher than the detonation velocity of the best conventional explosives, it is interesting to note that the energy density in the nuclear explosive is some 100 million times higher. The principle reasons for this disparity are that velocity only increases with the square root of energy density, and the high mass density of the fission explosive.

3.5.6 Thermal Waves with Hydrodynamic Flow

To review briefly, Section 3 initially treated the subjects of radiation transport and hydrodynamics separately:

- Section 3.3.4.2, Radiation Heat Conduction, discussed the process of heat transport by radiation in the absence of hydrodynamic flow;
- Section 3.4, Hydrodynamics, discussed shock hydrodynamics in the absence of radiation transport.

So far in this section we have discussed how shocks of increasing strength give rise to radiation transport as an important mechanism. Now I want to turn this around, and consider how radiation heating can give rise to hydrodynamic shocks. This is very important when we are dealing with the heating of cold matter by a radiation flux before a shock wave has become established, or in any situation where a radiation flux is the essential force in driving the hydrodynamic phenomena (especially if the flux varies with time).

As we have seen, the behavior of the driving temperature behind a thermal wave with respect to time determines the how the wave velocity varies with time. For a constant temperature radiation diffusion has a marked tendency to slow with time and penetration distance, decreasing as $t^{-(1/2)}$.

Now the pressure in the heated zone is the sum of the kinetic and radiation pressures (given by Eq. 3.1.7-10):

P_total = n*k*T + (4*sigma/3*c)*T^4

and the velocity of a compression wave produced by this pressure (from Eq. 3.5.5-1) is:

D = (((gamma + 1)*P_total)/(2*rho))^(1/2)

assuming as we have with other very strong shocks that the initial pressure is negligible compared to P_total, and the compression approaches its limiting value. This shows that for a constant temperature T_0, the compression wave velocity is also constant. If the process continues long enough, the wave must eventually overtake the thermal wave front and drive ahead of it.

This shock, that now leads the original thermal wave, may very likely be itself a radiating shock wave which is preceded by a thermal wave. But the difference is this - the new temperature driving this thermal wave (call it T_1) is different from the original temperature T_0 that created the driving pressure for the shock. The new T_1 is produced by the process of shock compression, and is likely much cooler than T_0. Also the location of the driving temperature source is no longer fixed with respect with to the medium of conduction, it now advances with the shock wave at velocity D.

Conversely, the process of radiation conduction behind the shock front has been greatly altered also. The original thermal wave is now penetrating in a moving medium, that has been compressed and preheated.

The pressure behind the leading shock remains in equilibrium, but the physical state of the medium is not uniform. Behind the point where the shock overtook the thermal wave, the pressure is created by a high temperature (T_0) at the original density of the medium. Ahead of this point, the density is higher, but the temperature is lower.

Now suppose that the T_0 is not constant, but is increasing with time. Depending on the rate of increase, the thermal wave may have a constant velocity, or be increasing as well. But this also creates increasing pressure, and an increasing shock velocity. Will the shock overtake the thermal wave now? Through some sophisticated analysis techniques, general laws have been developed that describe the conditions under which the pressure wave will overtake the thermal wave, or not.

Let the driving temperature vary by the power law:

Eq. 3.5.6-1

$$T_0 = \text{constant} * t^q$$

where q > 0 (if q=0 we have the case of constant temperature described above). From Eq. 3.3.4.2.2-14 (which shows the relationship between depth of penetration x, and chi and t) and Eq. 3.3.4.2.2-8 (which shows that chi ~ T^n)we get:

Eq. 3.5.6-2

$$x \sim (\chi t)^{1/2} \sim T^{n/2} * t^{1/2} \sim t^{(n*q + 1)/2}$$

where n is the opacity-related exponent for the medium of conduction. The speed of the thermal wave is:

Eq. 3.5.6-3

$$dx/dt \sim x/t \sim t^{(n*q - 1)/2}$$

If we assume that kinetic pressure dominates radiation pressure (as must be the case at least for the early phase whenever cold material is heated) then we can also say:

Eq. 3.5.6-4

$$D \sim T^{1/2} \sim t^{q/2}$$

Analysis shows (see Zel'dovich and Raizer, *Physics of Shock Waves and High-Temperature Hydrodynamic Phenomena*, pg. 677) that if:

$$(nq - 1)/2 < q/2$$

(showing that the shock velocity increases with time faster than thermal wave velocity), and

$$q < 1/(n - 1)$$

then the situation described above holds - the thermal wave initially runs ahead of the shock, but the shock eventually overtakes it, and leads the thermal wave for all later times.

The opposite situation exists if:

$$(nq - 1)/2 > q/2$$

(showing that the thermal wave velocity increases with time faster than shock velocity), and

$$q > 1/(n - 1).$$

In this case the shock initially leads, but is later overtaken by the thermal wave, after which the significance of the hydrodynamic motion steadily shrinks as the thermal wave moves farther and farther ahead (i.e. the ratio of shocked material to thermal wave heated material becomes small).

In the case that:

$$(nq - 1)/2 = q/2,$$

and

$$q = 1/(n - 1)$$

the speeds of the shock and thermal waves increase at the same rate. Separate phases of hydrodynamic an heat conduction transport do not exist, and the medium is heated and set in motion almost simultaneously.

3.6 Shock Waves in Non-Uniform Systems

So far I have discussed shock waves passing through a uniform medium where physical properties like density do not change. We are necessarily interested in the behavior of shocks in non-uniform systems since nuclear weapon design techniques include:

- Compression of one material (uranium say) by shock waves generated in another (high explosive); and
- Generation of shock waves by collisions between two materials.

3.6.1 Shock Waves at an Interface

The dimensions of the components of a nuclear weapon are small. This means that shock waves propagating in these components must eventually reach a boundary (in contrast explosion shocks in the atmosphere or ocean may die out without encountering a boundary). Since the most important physical property of a material for determining shock behavior is density, two basic types of interfaces exist:

- an interface where the shock enters a material of lower density; and
- an interface where the shock enters a material of higher density.

A shock wave is not affected by an interface ahead of the shock until the leading edge of the wave reaches it. In classical shocks this means that the shock propagates in a uniform manner until the shock front is actually at the interface at which point it will undergo an instantaneous change in behavior. In radiation shocks the leading edge is the edge of the Marshak wave preheating zone. This results in a more gradual interaction.

Two conditions that must be satisfied when a shock reaches an interface are:

- the pressures on both sides of the interface must equal; and
- the particle velocities on both sides of the interface must equal.

3.6.1.1 Release Waves

When shock waves meet a boundary with a lower density material a release wave is generated. This is a rarefaction wave, a reduction or release in pressure that converts internal energy of the shock compressed gas into kinetic energy.

The limiting case of a release wave is the arrival of a shock at a vacuum interface, also called a free surface. We can claim that the shock disappears in any meaningful sense at this point. All we have is a mass of compressed, moving material adjacent to a vacuum. The compressed material immediately begins to expand, converting the internal energy into even greater kinetic motion. In gases this expansion process is called escape. The expansion in turn creates a rarefaction wave that propagates back into the shock compressed material at the local speed of sound c_s.

3.6.1.1.1 Free Surface Release Waves in Gases

Let us assume that material supporting the shock wave is a gas. Since the pressure change in a rarefaction wave is always continuous (no instantaneous pressure drops), the pressure at the leading edge of the escaping gas is zero (in keeping with the requirement of equal pressures at the gas/vacuum interface) and the process of converting internal energy into kinetic energy is complete. The front edge thus immediately accelerates to the final maximum velocity (escape velocity) when it reaches the free surface. Farther back in the release wave the pressure increases and the velocity decreases. At the release wave boundary (which moves backward at c_s) the pressure and velocities are the same as in the original shock wave.

The escape velocity of a gas is equal to:

Eq. 3.6.1.1.1-1

$$u_{escape} = (2*c_s)/(gamma - 1)$$

If we use a frame of reference in which the unshocked gas was stationary we get:

Eq. 3.6.1.1.1-2

$$u_{escape} = (2*c_s)/(gamma - 1) + u_{particle}$$

and the rear edge of the release wave travels backwards at

Eq. 3.6.1.1.1-3

v_release = u_particle - c_s

The release wave thus stretches itself out with time at a speed of:

Eq. 3.6.1.1.1-4

((2/(gamma - 1)) + 1)*c_s

which for a perfect monatomic gas is 4*c_s.

3.6.1.1.2 Free Surface Release Waves in Solids

If the material is not a gas but a solid (and the shock has not vaporized it), the situation is a bit different. The solid cannot expand indefinitely, so expansion halts when it reaches the normal zero pressure density. The release propagates backward at c_s as before. In metals at moderate shock strengths (1 megabar and below), this expansion causes a phenomena called "velocity doubling", since the release wave doubles the velocity imparted by the shock. This doubling law begins to break down when shock strengths are high enough to cause substantial entropy increases.

Another important difference in solids is the existence of tensile strength. We have seen that in detonation waves that the pressure drops rapidly a short distance behind the shock front. When high explosives are used to generate shocks in metal plates a phenomenon called "spalling" can occur due to the interaction between the tensile properties of the metal, the release wave, and this shock pressure drop.

As the rarefaction wave moves back into the plate, it causes the pressure to drop from the shock peak pressure to zero. As it moves farther back into the plate it continues to cause pressure drops of the same absolute magnitude. When it encounters pressures below the shock peak this means that the pressure in the plate goes negative, that is, tensile stress is produced (you can think of this as the faster moving part of the plate pulling the slower moving part of the plate along). If the negative pressure exceeds the tensile strength of the metal, the plate will fracture. The fracture may simply open a void in the plate, but often the velocity doubled layer will peel off entirely and fly off the front of the plate.

If the original plate was actually a stack of two or more plates, the shock propagation into the plate stack would occur as before. However when the release wave reaches the boundary between plates, and if tensile stress exists at that point, the front plate will fly off and a new release wave will begin at the front side of the plate behind it.

3.6.1.1.3 Shock Waves at a Low Impedance Boundary

If the boundary is with a second material of lower density, rather than a true vacuum, we have a low impedance boundary. If this second material has a density *much* lower than the first, then the situation is essentially the same as the free surface case as far as the first material is concerned. The pressure at the release wave front will be negligible compared to the shock pressure and the escape velocity will be virtually the same.

The increased particle velocity at the front of the release wave acts like a piston in the second material, driving a shock wave ahead of itself. Since the particle velocity is steady behind a shock front, and the particle velocity on both sides of the material interface must be the same, we can see that the particle velocity behind the shock front in the second material will be the same as the velocity of the release wave front (i.e. the velocity of material ahead of a piston is the same as the piston itself). If we assume that the second material is a gas, and that the shock is strong enough to reach the limiting compression, then the second shock wave velocity will be:

Eq. 3.6.1.1.3-3

$$D_{2ndshock} = u_{escape} * (gamma + 1)/2$$

The pressure in this shock will be negligible compared to the original shock of course.

If the density of the second material is not negligible compared to the first, the result is intermediate between the case above and the case of a shock propagating through a material of unchanging density. The pressure at the interface drops, but not as drastically. The velocity of

the release wave front is lower, and creates a shock in the second material that is slower, but with a higher pressure. The rear of the release wave travels backwards at the same speed (c_s), but since the pressure drop at the release front is not as sharp, and it moves forward more slowly, the pressure gradient in the release wave is not as steep.

3.6.1.2 Shock Reflection

When shock waves meet a boundary with a higher density material a shock wave is reflected back into the first material, increasing its pressure and entropy, but decreasing its velocity.

3.6.1.2.1 Shock Waves at a Rigid Interface

The limiting case of an interface with a higher density material, is an infinitely rigid boundary. With such an interface there is no shock wave transmitted into the second material, all of the energy is reflected. This causes a doubling of pressure, and a decrease in particle velocity to zero. An interface between air and a solid (like the ground) is essentially a rigid interface.

3.6.1.2.2 Shock Waves at a High Impedance Boundary

If the second material is not enormously denser than the first, a transmitted shock is created in the second material, along with the reflected shock in the first. The shock reflection causes an increase in pressure and decrease in particle velocity in the first material. Continuity in pressure at the interface indicates that the transmitted shock must also be higher in pressure than the original shock, but the particle velocity (and shock velocity) will be lower.

3.6.2 Collisions of Moving Bodies

The detonation of high explosives is one way of generating shock waves. Another is high velocity collision between two materials.

When two bodies come into contact at a given velocity, the material at the colliding surface of each body must instantly change velocity so that the relative motion of the two surfaces is zero (due to symmetry we may consider either one to be stationary, or both to be moving, at our convenience, as long as the collision velocity is the same). A pressure must be generated by the collision exactly sufficient to accomplish this velocity transformation.

To illustrate the situation, first consider the head on collision between two moving plates of identical composition, and equal mass and velocity magnitude $|V|$. At the moment of contact the colliding surfaces will come to rest, and a shock wave will begin to propagate into each plate of sufficient strength to cause a velocity change of magnitude $|V|$. If $|V|$ is not too great, then the shock velocity will be close to the acoustic speed in the plates. When the shocks reach the opposite surface of each plate, all portions of both plates will be at rest. A release wave will begin at the plate free surfaces, resulting in velocity doubling. When the release waves reach the original impacting surfaces, both plates will have returned to their original (zero pressure) physical state, but they will now have velocities of magnitude $|V|$ headed in opposite directions to their original motion; they will have "bounced" off each other.

What we have just described is an elastic collision, a collision in which no kinetic energy is dissipated. This the limiting case for colliding bodies at low velocities. If the collision velocities are below the acoustic speed in the bodies then the collision forces are transmitted by acoustic waves. If the velocities become too large then true shock waves are formed. At still higher velocities, entropic heating and equation of state effects start to deviate from the velocity doubling rule. At this point the collisions become progressively less elastic with increasing speed.

Typically we encounter colliding bodies in physical systems in which we view one body as stationary. The impacting body travels at velocity V, and after collision shock waves propagate into the bodies with u_p equal to V/2 (if they are of identical composition). The actual shock front velocity depends on the degree of compression, which can be found by consulting the Rankine-Hugoniot relations, and the initial state and the EOS of the material.

3.6.2.1 Collisions of Bodies With Differing Impedance

The description above assumed that the colliding bodies had identical densities and compressibilities, which is often not true. When a low density body collides with a high density body, a higher velocity shock is transmitted into the lower density body. If a more high compressibility body collides with low compressibility body, then the shock-mediated velocity change of the high compressibility body will be much greater (the situation is identical to that described in 3.6.1.2.2 Shock Waves at a High Impedance Boundary, except that the high compressibility/low impedance material has not been previously shocked).

An interesting situation occurs if a low impedance material is trapped between two high impedance bodies during a collision. This may occur if one or both high impedance bodies have low impedance material covering their surface, or it they are not colliding in an effective vacuum and a stationary fluid is caught between them. For specificity, let us assume that a moving high impedance body collides with a stationary body consisting of a low impedance and high impedance layer.

At the moment of collision, a relatively low pressure, high velocity shock is set up in the low impedance material, which is greatly accelerated. The colliding body in decelerated to a much smaller extent. The shock is transmitted through the low impedance layer, and is reflected at the high impedance boundary, bringing the material to a near-halt again. The intensified reflected shock travels back to the first body, and is reflected again. Each reflection increases the pressure of the shock, and the density of the low impedance material, and incrementally transmits momentum from the first high impedance body to the second. After enough reflections have occurred the first will have been decelerated, and the second body accelerated sufficiently to bring them to rest with respect to each other. In the process, the low impedance material will have been highly compressed by repeated reflected shocks.

3.6.3 Collisions of Shock Waves

Shock reflections at an interface are an example of a shock wave collision. The general case of head on shock collisions presents no additional difficulties. From the point of collision, waves must propagate in either direction. The properties of each wave must be such that they bring the pressures and particle velocities on both sides of the collision interface into equilibrium. Symmetrical shock collisions act exactly like a shock wave encountering an infinitely rigid interface from the point of view of both shocks.

Entirely new phenomena occur though when shock collisions are not head on. The discussion below assumes we are dealing with strong shocks and/or dense matter (the effects are somewhat different for low density gases and moderate or weak shocks).

At very small angles, the point of collision moves very rapidly along the interface (V/Sin alpha; where alpha is the collision angle, 0 for head on). If the speed with which the collision point moves is supersonic with respect to the material behind the reflected shocks then the only new phenomena are that the shocks are reflected at an angle, and the pressure behind the reflected shocks is increased.

As alpha increases, the speed of travel of the collision point decreases and the reflected shock pressure increases. At a certain critical value of alpha (generally around 70 degrees) the collision point travels at sonic velocity or less, and a new shock front is formed by the collision (a phenomenon called the Mach effect). This shock, called the Mach stem, bridges the two principal colliding shocks, replacing the V-shaped two-shock collision with a three sided shock collision. The stem travels parallel to the direction of travel of the collision point.

At the critical angle where the Mach front is formed the pressure reaches to a maximum, which can be some six to eight times the original shock pressure. Above the critical angle, the pressure drops again. When alpha reaches 180 degrees the two shocks are no longer colliding and the pressure drops back to its original value.

3.6.4 Oblique Collisions of Moving Bodies

A phenomenon of some importance related to the Mach effect arises when oblique collisions are between two bodies in a void, not shocks travelling through a continuous medium. In this case there is nothing for a shock to travel through ahead of the collision point, so a Mach stem cannot form. Instead at the critical angle where the collision point is sonic or subsonic with respect to the reflected shock, a high velocity jet of material is ejected along the collision point trajectory. This is variously called the Neumann, Monroe, hollow charge, or shaped charge effect. When the collision is cylindrical or conical instead of planar, a jet is always created regardless of the angle of collision.

This effect is used in conventional weapons systems such as anti-tank missiles to create jets capable of penetrating thick armor. The effect can also arise in implosion systems when shock wave irregularities occur. In this case it is very undesirable because it disrupts the implosion.

3.7 Principles of Implosion

The term "implosion", which denotes a violent inward collapse or compression, fits many different physical schemes for rapidly compressing materials to high densities. An implosion may be adiabatic or shock induced, or both; the geometry of the compression may be one, two, or three dimensional. This subsection surveys these various implosive systems, describing their principles and properties.

3.7.1 Implosion Geometries

An compression process can be symmetric about one, two, or three spatial axes.

A good example of one dimensional compression (linear compression) is the compression of the fuel/air mixture in the cylinder of an internal combustion engine. If r_0 is the original length of the gas column in the cylinder, and r_1 is the length after compression, then the density increase is:

Eq. 3.7.1-1

$$\rho_1/\rho_0 = r_0/r_1$$

That is to say, it is inversely proportional to the change in scale (the relative change in length).

Two dimensional compression (cylindrical compression) can be thought of as squeezing a tube so that its radius decreases uniformly (it doesn't get squashed flat). If r_0 denotes the original radius, and r_1 the radius after compression then we can say:

Eq. 3.7.1-2

$$\rho_1/\rho_0 = (r_0/r_1)^2$$

I.e. it is inversely proportional to the SQUARE of the change in scale.

Three dimensional compression (spherical compression) can be thought of as squeezing a sphere so that its radius decreases uniformly. In this case we can say:

Eq. 3.7.1-3

$$\rho_1/\rho_0 = (r_0/r_1)^3$$

I.e. it is inversely proportional to the **CUBE** of the change in scale.

For the same change in scale, a higher dimensional implosion produces a much greater degree of compression. The relatively sluggish linear case in fact is rarely though of as being an "implosion".

The spherical implosion gives the most rapid compression and, being symmetrical in all directions, is also relatively easy to analyze theoretically. It is the most widely used geometry for implosion in nuclear weapon designs.

3.7.2 Classes of Implosive Processes

In addition to geometry, another critical factor in characterizing implosion systems is the pressure-time curve that actually causes the compression.

At one extreme, we have gradual homogenous adiabatic compression. In this case the pressure exerted on the implosion system increases continuously, at a slow enough rate that the pressure within the implosion system is uniform everywhere. This type of implosion does not increase entropy and allows arbitrarily high degrees of compression, if sufficient force can be exerted.

The other extreme is shock compression. Here the pressure increase is instantaneous, and entropy is always increased. There is a theoretical limit to the degree of compression achievable in this manner, regardless of the shock pressures available.

Other important cases can be defined that lie between these extremes, depending of the pressure-time function. For example, pressure curves that increase exponentially with time (physically unachievable, generally speaking) allow infinitely high compression in arbitrarily short length of time.

In characterizing the properties of different classes of implosions we use spherical compression as the model, due to its theoretical simplicity, and practical interest.

The reader interested in exploring classes of implosions and their properties is referred to the excellent papers by J. Meyer-ter-Vehn (and others):

Meyer-ter-Vehn, J.; Schalk, C. 1982. *Selfsimilar Spherical Compression Waves in Gas Dynamics*, Z. Naturforsch. 37a, 955-969.

Meyer-ter-Vehn, J. 1992. *Physics of Inertial Fusion*, in Proceedings of the International School of Physics "Enrico Fermi", Course CXVI, *Status and Perspectives on Nuclear Energy: Fission and Fusion*, 395-423.

3.7.3 Convergent Shock Waves

This is the most obvious method of achieving implosion, and was the first two and three dimensional implosion system to be developed (during the Manhattan Project). The idea is to create a cylindrical or spherical inward directed shock wave which converges on the material to be compressed.

In a convergent shock wave the flow of matter behind the shock front is also convergent (inward directed). This means that following the shock compression there is an adiabatic compression phase in which the kinetic energy of the shocked material is converted into internal energy, causing the flow to decelerate. The pressure behind the shock front, rather than remaining constant as in a classical shock wave, continues to rise. This increasing pressure is transmitted to the shock front, causing it to continuously strengthen and accelerate.

In theory the acceleration of the shock front continues without limit as the shock approaches the center, achieving infinite velocity, pressure, and temperature. Since the area of the shock front is diminishing, the actual total energy in the shock front decreases, reaching zero at the center.

Although the increase in shock strength is unbounded in principle, subject only to the granularity of atomic structure, in practice the symmetry of the implosion breaks down long before then, providing a much lower limit to the maximum degree of energy cumulation. In contrast to plane or divergent shocks, which are stable, convergent shocks are unstable and irregularities will grow with time.

The increase in density at the shock front is bounded of course, and so is the maximum density due to adiabatic compression. In fact, when the shock reaches the center, the density is everywhere uniform. The temperature and pressure profile established at this point increases without limit toward the center.

The degree of shock intensification depends on the geometry (planar, cylindrical, or spherical), and the equation of state of the material. It is higher for less compressible materials.

The scaling laws make use of the parameter alpha given in the table below (assuming a gas law EOS):

Table 3.7.3-1. Alpha Values for Varying Gammas and Geometries

Gammas:	1	6/5	7/5	5/3	3
Cylindrical	1	0.861	0.835	0.816	0.775
Spherical	1	0.757	0.717	0.688	0.638

The velocity of the converging shock front is given by:

Eq. 3.7.3-1

$$V = K1*r^{[(alpha-1)/alpha]}$$

where alpha is selected from the above table, and the proportionality constant K1 is determined by initial conditions (V_0 and r_0). If the value of gamma is 3 (close to that of uranium) in a spherical implosion then alpha is 0.638. This gives us:

Eq. 3.7.3-2

$$V_{Usphere} = K1*r^{(-0.567)}$$

The shock velocity is thus approximately proportional to the inverse square root of the radius.

The pressure behind the shock front is given by:

Eq. 3.7.3-3

$$P = K2*r^{[2*(alpha-1)/alpha]}$$

Again K2 is established by initial conditions (p_0, r_0). For gamma=3, this is:

Eq. 3.7.3-4

P_Usphere = K2*r^(-1.134)

The pressure is thus approximately proportional to the inverse of the radius.

Side note:

The scaling laws for implosion were obtained by consulting several mutually supporting sources that analyze the theoretical properties of converging shock waves, which cross checks with published data on converging shock waves. These laws also appear generally consistent with the known properties of the Fat Man implosion bomb.

There is a curious belief found in the public domain nuclear weapons literature that implosion shock velocity scales with the inverse square (**NOT** square root) of the radius, and thus the pressure scales as the inverse FOURTH power of the radius (see for example Lovins, 1980, Nature V. 283, pg. 821). Since the area of the shock front is proportional to r^2, and the energy density in a shock wave is proportional to the pressure, this belief leads to the astonishing (and impossible) conclusion that the TOTAL energy in the shock increases rapidly, and without limit (i.e. it is a source of "free energy", in contravention to the laws of thermodynamics).

The origin of this belief appears to be a comment in a biography of Edward Teller (*Energy and Conflict: The Life and Times of Edward Teller*), judging by a statement made by De Volpi in *Proliferation, Plutonium, and Policy*, 1979, pg. 300.

The time for the shock wave to reach the center from r_0, and V_0 is

Eq. 3.7.3-4

$$t_c = (r_0/V_0)*alpha$$

The location of the shock front at time t is:

Eq. 3.7.3-5

$$r(t) = (t_c - t)^alpha * K3$$

Where K3 can be calculated by:

Eq. 3.7.3-6

$$K3 = (r_0^(1-alpha) * V_0^alpha)/(alpha^alpha)$$

The compression achieved at the shock front has the classical shock compression limit: nu = (gamma + 1)/(gamma - 1).

Now the maximum degree of compression would be achieved if all of the kinetic energy were converted into internal energy, bringing the imploding gas to rest. We can make a somewhat naive estimate of what this compression limit is by remembering that in an extremely strong shock, energy is divided equally between kinetic energy and internal energy. Then by using the gas law EOS and the adiabatic compression law, we can calculate the additional density increase if the kinetic energy were converted to internal energy adiabatically:

Eq. 3.7.3-7

$$rho_2/rho_1 = 2^(1/gamma-1)$$

where rho_1 is the shock compressed density and rho_2 is the adiabatically compressed density.

This gives us for gamma = 5/3 a density increase of

Eq. 3.7.3-8

rho_2/rho_1 = 2^1.5 = 2.82, or rho_2/rho_0 = 11.3.

For gamma = 3 (more typical of condensed matter) the increase is:

Eq. 3.7.3-9

rho_2/rho_1 = 2^0.5 = 1.41, or rho_2/rho_0 = 2.82.

This estimate is naive because the state of zero flow is never achieved. When the shock reaches the center, the gas is still flowing inward everywhere. An accurate value for rho_2/rho_0 at this moment, when gamma = 5/3, is 9.47.

In an extremely intense implosion shock that ionizes the material through which it passes the initial shock compression is much greater than in an ideal gas. The shocked gas acts as an ideal monatomic gas with a gamma of 5/3. If we assume an initial compression of 12 (3 times higher than for an ideal monatomic gas) then the final density after convergence will be on the order of 30.

3.7.3.1 Convergent Shocks With Reflection

When a converging shock reaches the center of the system it should, theoretically, create a reflected shock moving outward driven by the increasing pressure toward the center. This shock is a relatively slow one (compared to the imploding phase). The reflected shock front reverses the inflow of gas. Behind the front the gas is expanding outward, again quite slowly compared to its inward velocity. The density reaches a new maximum at the shock front, but falls rapidly behind it reaching zero at the center. This is a consequence of the increasing entropy toward

the center, but is augmented by the fact that pressure drops behind a diverging shock front. The new density maximum is thus a local phenomenon at the shock front. Matter is actually being driven outward from the center and for any volume around the center, the density falls continuously once the reflected shock front has passed. You can think of the implosion as having "bounced" at the center, and is now rebounding outward.

The compression achieved at the shock front for gamma = 5/3 in spherical geometry is rho_3/rho_0 = 32.0 For 7/5 it reaches a value of 216. For cylindrical geometry and gamma = 5/3, rho_3/rho_0 is 22.8. If we get a factor of 4 initial compression enhancement due to an ionizing shock, and a gamma of 5/3 afterward, then the compression for spherical geometry will be 96, and 68.4 for cylindrical geometry.

Whether the high densities achieved at the outgoing shock front are "useful" or not depends on the characteristic length scale of the physical processes being enhanced by the compression. In a fission reaction the scale is given by the neutron mean free path, the average distance a neutron travels between interactions, which is typically a few centimeters. If the thickness of the shock compressed layer is less than this, the reaction will be more affected by the average density of the core than by the local density at the front. In a fusion reaction the reaction rate at a given temperature is governed by the local density of the fuel, so the reaction rate will be enhanced by the rebounding shock.

It appears that the peak compression at the reflected shock is less than that predicted by theory and it not hard to guess why. Before this second symmetric shock phase begins the implosion has to pass through a singularity in the center where the shock front decreases to zero size, and increases to infinite strength. Since implosion symmetry breaks down before this point, we actually get a turbulent flow that dissipates some of the inflowing kinetic energy as heat. We thus get an expanding stagnation front with higher entropy increases, and lower compressions than theory would indicate.

The problem of non-ideal behavior near the center can be alleviated by placing a high impedance boundary around the center of the system - a small high density object - to reflect the shock back outward before it converges to the point where asymmetry becomes severe. This also reduces the amount of energy deposited as heat at the center, which is probably an advantage, but it starts to change the character of the implosion from the simple ideal shock implosion model.

If the main imploding mass is surrounded by a higher density imploding shell (which presumably originated the initial convergent shock) then additional compression is achieved when the outgoing shock is reflected at this high density interface. If a high density central mass is also present a series of alternating inward and outward reflected shocks can be created that achieve very high densities. In this case it is the kinetic energy of the decelerating high density layer (acting as a pusher/tamper) being brought to rest that is supplying the additional compressive work. Such a system is of course much more complex than a simple convergent shock implosion.

3.7.4 Collapsing Shells

A variation of implosion that is of considerable interest is the collapse of a hollow cylindrical or spherical shell. We consider two cases here:

1. a shell initially at a uniform velocity, but not acted on by outside forces;
2. a shell initially at rest, but reacting to an external pressure.

It is apparent that if another body is located at the center of the hollow cavity, a collapsing shell can be used to create a convergent shock with reflection.

3.7.4.1 Shell in Free Fall

The first case is often referred to as a "free falling" shell, although it is not really falling, just imploding due to inertia. We assume that every particle in the shell has a velocity vector directed towards the center of the hollow shell, and that the magnitude is uniform everywhere.

This situation is typical of a thin shell that has just been accelerated by an imploding shock wave after the velocity doubling pressure release has been completed.

As the shell implodes the area of its outer surface gets smaller. This means that the shell is either getting thicker, or it is getting denser, or both.

If it is getting thicker, then the inner surface is moving farther than the outer surface, which means that it is also moving faster than the outer surface. Since they were initially moving at the same speed, the inner surface must be accelerating. Due to the conservation of momentum and energy, this means that the outer surface is also slowing down. This is the only situation possible in a shell made of incompressible material.

This transfer of momentum creates a pressure gradient in the shell. Moving from the outer surface inward, the pressure increases from zero at the surface to a maximum that marks the division between the accelerating and decelerating regions. The pressure then drops to zero again at the inner surface. This maximum pressure will be much closer to the inner surface than the outer surface. Real materials are compressible to some extent so the pressure increase inside the shell is also accompanied by an increase in density.

3.7.4.2 Shell Collapse Under Constant Pressure

To investigate the second case, let us assume that the shell is initially at rest with no forces acting upon it. Then let a uniform external pressure be applied suddenly (we will ignore the mechanism by which the pressure is generated). The sudden pressure jump will create an acoustic of shock compression wave that will move inward, accelerating and compressing the shell. When the wave reaches the inner surface, a release wave will be generated that will move outward at the speed of sound. When the edge of the release wave, the point where the pressure drop begins, reaches the outer surface it will halt (we assume it cannot propagate beyond the shell, or alter the pressure being exerted). A pressure profile will have been

established in the shell, dropping from the maximum applied pressure at the outer surface to zero at the inner surface. Similarly, a continuous density and velocity profile will exist.

If the shock was strong enough to convert the shell into a gas, then the inner surface will begin imploding at escape velocity when the shock reaches it. The leading edge of the implosion, travelling at high speed, but with zero density and pressure will reach the center of the void far ahead of the outer surface. If the shell remains solid, then the inner surface will undergo "velocity doubling", and drop to its zero pressure density.

The external pressure, not balanced by any internal pressure until the inner surface reaches the center, will cause the shell to continue to accelerate. The acceleration will be transmitted to the inner surface by the pressure gradient.

3.7.5 Methods for Extreme Compression

What densities are actually achievable with real implosion systems and real materials? We must consider actual compressibility limitations, and how close physical implosion systems can get to the ideal limiting case.

Although the value of gamma for condensed matter at zero pressure (and compression) is greater than the 5/3 of a perfect gas, this is partially offset by the fact that the value of gamma always falls as the density of condensed matter increases (approaching 5/3 as a limit). However, extremely high shock pressures are required to approach the limiting shock compression value, even using the STP gammas. For example the STP value of gamma for uranium (2.9) drops to 2.4 when its density has been doubled. Using the STP value we can estimate a limiting compression for uranium of ((2.9 + 1)/(2.9 - 1) = 2.05.

This degree of compression is actually achieved at a shock strength of 7 megabars, approaching the practical limits of high explosive implosion. A static pressure of 5 megabars would have been sufficient to achieve this density through adiabatic compression. The 2 megabars of "lost compression" in the shock is due to energy expended in heating the uranium. The inefficiency of shock compression starts to increase rapidly above this point. A 20 megabar shock compresses uranium by a factor of 2.5, comparable to 12 megabars for shockless compression. In extremely strong shock waves, the energy of the shock is divided equally between heating and kinetic energy, with a negligible portion going into actual compression.

To achieve much higher compressions we have two options: successive shock waves, or adiabatic (i.e. non-shock induced) compression.

The description of a single shock implosion given in subsection 3.7.2 shows that adiabatic compression naturally follows the implosion shock, but the amount of compression achievable is still limited. If much higher densities are required some method of enhancing the adiabatic compression process is required.

By dividing the pressure increase into two or more shock steps, the inefficiency of very strong shocks can be dramatically reduced. In fact, two shocks are sufficient to eliminate most of the entropic heating for pressure increases up to many tens of megabars.

If the pressure increase is divided between a large number of weak shocks, the effect is essentially identical to adiabatic compression. In fact, true adiabatic compression can be considered the limiting case of an infinite number of infinitely weak shocks. Conversely in very rapid adiabatic compression where the pressure gradients are steep, the inherent instability of compression waves tends to cause the pressure gradient to break up into multiple shocks. Thus for practical purposes the multi-shock and adiabatic compression approaches are closely related, and either model can suffice for analyzing and modelling extreme compression.

Two general styles of multi-shock implosion are possible. The first has already been described - a convergent shock system with reflection. In this system, each successive shock is going in the opposite direction as the previous one, and requires shock drivers on each side. An alternative approach is unidirectional shock sequences. In this approach a succession of shocks travel in one direction in "follow-the-leader" style.

In an optimal unidirectional multi-shock implosion system the shocks would be timed so that they all reach the center simultaneously (each successive shock is faster and will tend to overtake the earlier shocks). If they overtake each other earlier, then a single intense shock will form. If they do not merge before reaching the center, then shock wave rebounds will collide with incoming shocks, creating complicated shock reflection effects. Similarly the time pressure curve in adiabatic compression should avoid steepening into a single strong shock, or breaking up into a rebounding shock sequence.

The maximum speed of high density compression is determined by the speed of the leading edge of the pressure gradient, and the thickness of the material to be compressed. In true adiabatic (shockless) compression, this speed is the acoustic velocity in the uncompressed material. Since normal acoustic velocities are on the order of 1-3 km/sec, and material thicknesses are at least a few cm, this compression time is on the order of 10-100 microseconds. And since nuclear bombs typically disassemble in a microsecond or less, it is clear that an initial strong shock with a velocity approaching 100 km/sec is unavoidable.

Let us consider an example where the fuel being compressed is Li6D, and the total pressure increment is 10,000 megabars. The density of the fuel being compressed (designated rho_init) is 0.82, and if the initial shock pressure (p_0)is 50 megabars then the shock velocity is roughly:

Eq. 3.7.5-1

$$D_0 = [p_0(\gamma + 1)/2*\rho_{init}]^{0.5}$$

= [((5 x 10^7 bars)*(10^6 dynes/cm^2/bar)*(1 + 5/3))/(2 * 0.82)]^0.5

= 9.0 x 10^6 cm/sec

Since this an ionizing shock in Li6D initially at STP, the density after initial shock compression (rho_0) may be as high as 12 or so.

From the adiabatic law:

P_1 * V_1^gamma = P_final * V_final^gamma

we get the following expression for the adiabatic compression ratio:

Eq. 3.7.5-2

(P_final/P_0)^(1/gamma) = V_0/V_final = rho_final/rho_0

If all of the remaining pressure increase is adiabatic and gamma = 5/3, then

rho_final/rho_0 = 200^(3/5) = 24.0

giving a final density of around 288.

The next question is how many discrete shocks are necessary to approach this performance. To examine this, we will assume a plane geometry for the shocks (i.e. they neither diverge or converge).

The compression ratio from a shock is:

Eq. 3.7.5-3

$$V_i/V_{i+1} = [(\gamma + 1)P_{i+1} + (\gamma - 1)P_i] / [(\gamma - 1)P_{i+1} + (\gamma + 1)P_i]$$

where i is the pre-shock state and i+1 is the post-shock state. If we divide the pressure jump from P_0 to P_{final} into n jumps with the same pressure ratio (P_r) then the ratio for a single jump is:

Eq. 3.7.5-4

$$P_r = P_{i+1}/P_i = \exp[\ln(P_{final}/P_1) / n]$$

where i = 1, 2 ... n (this numbering explains why the initial shock was called the zeroth shock).

The compression ratio with each shock in the sequence is:

Eq. 3.7.5-5

$$\rho_r = \rho_{i+1}/\rho_i = [(\gamma + 1)P_r + (\gamma - 1)] / [(\gamma - 1)P_r + (\gamma + 1)]$$

and the total compression by this sequence is obviously:

Eq. 3.7.5-6

$$\rho_{final}/\rho_0 = (\rho_{i+1}/\rho_i)^n = (\rho_r)^n$$

By computing various values of n, and assuming gamma = 5/3, for the earlier example we find:

Table 3.7.5-1. Total Compression Vs. Number of Shock Steps

n	Density Ratio (ρ_{final}/ρ_1)	Pressure Ratio (P_r)
1	3.92	200.00
2	10.07	14.14
3	15.20	5.85
4	18.26	3.76
5	20.05	2.89
10	22.92	1.70
25	23.84	1.24
50	23.98	1.11
infinity	24.00	1.00

We see from the table that on the order of five shocks can capture a large majority of the adiabatic compression. Another way of looking at it is: we are assured of losing little

compression if a continuous pressure gradient breaks up into shocks each with a pressure jump of less than 3 or so.

In this analysis I have implicitly assumed that the shock sequence is unidirectional (i.e. the compression is caused by a sequence of shocks travelling in the same direction). This need not necessarily be the case. The implications for compression efficiency in the reflected shock described system at the end of Section 7.2.3.1 are of course the same.

In a unidirectional shock sequence, to achieve maximum compression throughout the layer each shock should be timed so that they all reach the opposite interface of the material simultaneously. If the thickness of the material is d, then the transit time for the initial shock is:

Eq. 3.7.5-7

$$t_{transit_0} = d/D_0$$

Each shock in the following sequence travels through material that has been compressed and accelerated. For the ith shock we can say:

Eq. 3.7.5-8

$$D_i = [[(\gamma + 1)*(P_r^i)*P_0 + (\gamma - 1)*(P_r^{(i-1)})*P_0] / [2*(\rho_r^{(i-1)})*\rho_0]] + (1 - (1/\rho_r))*D_{i-1}$$

except for the first shock in the sequence where we have:

Eq. 3.7.5-9

$$D_1 = [[(\gamma + 1)*P_0 + (\gamma - 1)*P_0] / 2*\rho_0] +$$

(1 - ((gamma-1)/(gamma+1)))*D_0

Since the driving interface of the layer is accelerated by the initial shock, the distance required to traverse the layer is less for subsequent shocks. When combined with the increasing shock velocities in the initial frame of reference, and multiplicative nature of the pressure increases, we can easily see that the driving pressure must rise very rapidly close to t_transit_0.

Brueckner and Jorna [*Laser Driven Fusion*, Reviews of Modern Physics, Vol. 46, No. 2; April 1974; pg. 350] derive an expression for how pressure must change with time using an assumption of adiabatic compression, treating i as a continuous variable, taking gamma as 5/3, and apparently also treating d as being constant:

Eq. 3.7.5-10

P(t_i)/P_0 = [[3*(t_i/(t_transit_0 - t_i)) / (8*(5^0.5))] + 1]^5

Plotting this function of P(t_i) vs. t_i shows this quasi-asymptotic rise.

The geometry of the systems we are actually interested are not really planar, but spherical or cylindrical instead. Still, this model provides valuable insight. Since the fuel volume is proportional to r^3 (for spherical geometry) or r^2 (for cylindrical geometry), it follows that most of the fuel mass resides at a fairly small radial distance from the surface. Half of the fuel is within 0.103 radii of the surface of a spherical mass, 0.146 radii of a cylindrical one. The effects of shock front convergence at these small radial distances is slight.

The fact that the opposite side of a fuel layer containing most of the fuel is not a free surface also relaxes the pressure history requirements. It does not matter whether the shock sequence or pressure gradient converges to a single shock at this point since the shocks will continue onward toward the center. The pressure increase does not need to be as fast or steep. It *is* important to avoid the buildup of excessively strong shocks in this outer layer of fuel, but farther in toward the center it does not matter much since inefficient compression here has little effect on the overall compression.

How to achieve the desired pressure-time curve in a fusion bomb where the shock production system (the fission trigger) destroys itself in a very short burst of energy, is a difficult and intriguing problem.

Possible approaches include using:

1. a radiation barrier between the primary and secondary with variable opacity to create a suitable time-dependent energy flux;
2. a pusher that transfers momentum and energy to the fuel over time by multiple reflected shocks and adiabatic compression;
3. layers of materials of different densities (shock buffers) to create multiple reflected shocks (this option can be considered a variant of the second); or
4. modulating the production of energy to provide the desired quasi- exponential pressure-time curve directly (e.g. tailoring the neutron multiplication rate, use of fusion boosting, use of multiple staging, or some combination of these).

Finally, we should not over look the importance of ionization compression (Section 3.2.4 Matter At High Pressures). When matter is heated until it becomes highly ionized, dense high-Z material will tend to expand and compress low-Z material due to the particle pressure of the ionized electrons. As an example the initial densities of uranium and lithium deuteride are 0.058 and 0.095 moles/cm^3 respectively. If adjacent bodies of uranium and lithium deuteride are heated until the uranium reaches its 80th ionization state (81 particles per atom, counting the ion), the lithium and deuterium will be completely ionized, and the particle densities become

1. Uranium: 0.058 * (80 + 1) = 4.7
2. Lithium Deuteride: 0.095 * ((3 + 1) + (1 + 1)) = 0.57

a particle density ratio of 8.2. The uranium will tend to expand, compressing the lithium deuteride, until the particle densities are the same. Ionization compression can thus lead to compression factors above 8 in this case.

3.8 Instability

Instabilities arise when force is exerted (or work is done) at a fluid interface. In unstable situations small deviations from symmetry or planarity tend to grow with time, magnifying the original imperfections. Since in the real world perfect symmetry or plane surfaces never exist (if only due to the atomic granularity of matter), unstable situations inevitably disrupt themselves with time. The key question is whether the initial irregularity magnitude, the rate of growth, and the length of time the unstable situation lasts are together sufficient to cause significant disruption.

3.8.1 Rayleigh-Taylor Instability

Probably the best known instability is the Rayleigh-Taylor instability. It arises when a denser fluid exerts force on a less dense one. The classic example of this is trying to float water on oil. Even if a (nearly) perfect layer of water can be initially established on top of the oil layer, it quickly sinks through to the bottom. The water-on-oil system is not stable.

The water sinks to the bottom because, as is generally true in physical systems, the water-on-oil system evolves to a configuration of lower potential energy. The energy gained by the oil rising is more than offset by the energy given up by the water sinking. No matter how perfect the water layer is initially, it will eventually sink to the bottom unless other stabilizing forces exist (like surface tension).

The smoother the interface is, the slower the instability develops. By the same token it is a rapidly accelerating process. Each increase in irregularity leads to even more rapid growth. The growth in irregularities is in fact exponential with time. Even absolutely perfect interfaces cannot prevent this from occurring since Brownian motion will create atomic scale irregularities where none initially existed.

This issue arose during the Manhattan Project when the forces involved in atomic bomb assembly and disassembly were being considered. Atomic bombs are made of materials of different densities. Large forces are generated by high explosive implosion, and immensely larger ones are generated by the fission energy release. Do instability problems arise that constrain the design, or limit the efficiency of the bomb?

A typical law governing the growth of Rayleigh-Taylor instabilities is:

Eq. 3.8.1-1

$$P = e^{t*[a*k*(rho_1 - rho_2)/(rho_1 + rho_2)]^0.5}$$

where P is the perturbation amplification, a is acceleration, rho_1 and rho_2 are the densities of the denser and lighter liquids, t is time, k is the perturbation wave number (determined by the scale of the perturbation relative to the scale of the whole interface), and e is the natural logarithm base. From this we can see that the exponential growth rate is affected by the relative and absolute densities of the fluids, and the acceleration (that is to say, the force exerted across the interface).

It is evident from the above equation that for a fixed velocity increment, the amplification of irregularities is lower for higher accelerations (there is less time for growth). In the limiting case of shock acceleration, it does not occur at all. It is only when unbalanced forces exist for significant time intervals that the instability becomes a problem.

This can occur during adiabatic compression, the collapse of thick shells, or when expansive forces are being generated by explosive reactions. The possible instability of the fissile core/tamper interface in both Manhattan Project bomb designs caused substantial concern (later shown to be unwarranted).

Despite considerable attention paid to this potential problem, it appears that it is largely a non-issue in weapons design. Due to the extremely short time scales involved, and the dominance of shock compression phenomena, significant instability problems do not appear to arise often. In fission bombs this is mostly a problem during the adiabatic compression phase behind the convergent shock wave of an implosion. It can also be an important factor during the confinement and expansion of the fusion fuel in a thermonuclear secondary, where a low density fuel first halts the pusher implosion, then drives the pusher outward.

3.8.2 Richtmyer-Meshkov Instability

Unlike Rayleigh-Taylor, which does not arise in true shock compression, Richtmyer-Meshkov instability is a phenomenon only found during shock compression (or a very steep compression wave that approximates a true shock). This instability is caused by irregularities in an interface between two media of differing density, and the effects of shock velocity variation with density. When the shock arrives at the interface, the portions of the transmitted shock front in the denser medium will lag behind the portions in the less dense one. If the shock is initially travelling in a less dense medium, there will be a reflected shock as well. The reflected shock will also be irregular since the first part of the shock to encounter the interface will be reflected back before the later parts. An initially flat shock thus gives rise to transmitted, and possibly reflected shocks, that are no longer flat, and this waviness may tend to grow with time. In a plane shock, once the shock has finished crossing the interface, or has finished being reflected, the shock will tend to return to planarity due to the inherent stability of plane shock waves. Convergent shocks are not stable though, so Richtmyer-Meshkov can be a serious problem in implosion systems. In addition, the shock front irregularities generate vortexes through the Helmholtz instability effect behind the front which can lead to mixing that continues long after the shock has passed by.

Richtmyer-Meshkov instability also differs from Rayleigh-Taylor in that it will arise no matter which side of the interface the compression force arrives from. Shocks travelling from low density regions to high density ones, and from high density regions to low density both cause this phenomenon (although the later case causes transmitted shocks only). The growth of the instability with time is linear instead of exponential though, which tends to limit its severity.

3.8.3 Helmholtz Instability

This instability is found in systems where flow occurs parallel to the interface between two fluids (this sliding motion is also called "shear"). Slight wavy irregularities in the interface cause the flow on either side to be diverted form a straight flow into a longer path, just as air is diverted as it flows over a curved airplane wing. By Bernoulli's principle the pressure in this region must drop, as it does to create lift with a wing. These pressure imbalances cause the wavy irregularities to grow in amplitude. When the height of irregularities reach a certain fraction of their wavelength they will curl, just like a breaking ocean wave. The rotational motion of the curl quickly becomes a vortex, which mixes the two fluids together. The effect of Helmholtz instability is to create a mixing layer consisting of a series of vortices maintained and driven like roller bearings by the fluid flow on either side.

Helmholtz instability arises in implosion systems as a side effect of other processes. Mach reflection from a surface creates a region of fluid flow behind the Mach front(the slip stream), which is subject to Helmholtz instability. Shock reflection and refraction at an irregular interface associated with Richtmyer-Meshkov instability also produces regions of fluid flow that can cause Helmholtz instability effects.

4.0 Engineering and Design of Nuclear Weapons

This section collects material in the open literature to provide a coherent survey of nuclear weapons technology. Everything in this section is in the public domain (not the same thing as being unclassified however), or is reasonable extrapolation or speculation based on public domain material. It was prepared without my ever having had access to any material on nuclear weapons not in the public domain. Much of the source material has been in the public domain for decades. Lest I be accused of providing instructions for building atomic bombs, I will point out that the discussion is general in nature and nowhere are detailed designs described. The actual construction of even a simple device requires specific dimensions, masses, and specifications of composition. I do not provide these, nor have I ever seen or attempted to prepare such a detailed design.

Preparing an actual weapon design (without extensive experimentation with real explosive and nuclear materials) requires significant amounts of numeric modelling of hydrodynamic and neutron transport effects. I do not discuss these computational techniques at all, although this is mainly to avoid excessive technical detail since the methods themselves are also not classified and are readily available in standard texts.

To use an architectural metaphor, the type of information set forth here can be compared to a general description of building construction techniques. To actually construct a building however, detailed blueprints prepared by an architect are required. A survey of construction techniques merely gives one an idea of what kinds of buildings can be built, and allows one to make general estimates of things like the amount and kind of materials required and the cost. It does not provide one with the knowledge required to actually build one.

Lack of knowledge has never been an obstacle to any nation in developing nuclear weapons. The problem is in obtaining the necessary tools and materials. Nothing I say here could be of any interest or assistance to a nation pursuing a nuclear weapons program. Admittedly, the situation with respect to terrorist acquisition of weapons is a bit different. But even here, actual weapon construction requires a type of information I do not provide, and most importantly it requires access to the proper materials. Preventing access to these materials is the only way of providing security from the spread nuclear weapons. Suppressing discussion of unclassified or public domain information does not provide any measure of security. It provides instead only a potentially dangerous illusion of security.

Interestingly enough, the United States government conducted a controlled experiment called the *Nth Country Experiment* to see how much effort was actually required to develop a viable fission weapon design starting from nothing. In this experiment, which ended on 10 April 1967, three newly graduated physics students were given the task of developing a detailed weapon design using only public domain information. The project reached a successful conclusion, that is, they did develop a viable design (detailed in the classified report UCRL-50248) after expending only three man-years of effort over two and a half calendar years. In the years since, much more information has entered the public domain so that the level of effort required has obviously dropped further.

This experiment established an upper limit on the required level of effort that is so low that the hope that lack of information may provide even a small degree of protection from proliferation is clearly a futile one.

The material in this section has been shown to persons knowledgeable in the field, who have agreed with my assessment that this material does not constitute a proliferation risk. I have also offered to submit the material to the Department of Energy for review to determine whether any material constituting such a risk is present, but this offer has apparently been declined since I have received no response.

4.1 Elements of Fission Weapon Design

4.1.1 Dimensional and Temporal Scale Factors

In Section 2 the properties of fission chain reactions were described using two simplified mathematical models: the discrete step chain reaction, and the more accurate continuous chain reaction model. A more detailed discussion of fission weapon design is aided by introducing more carefully defined means of quantifying the dimensions and time scales involved in fission explosions. These scale factors make it easier to analyze time-dependent neutron multiplication in systems of varying composition and geometry.

These scale factors are based on an elaboration of the continuous chain reaction model. It uses the concept of the "average neutron collision" which combines the scattering, fission, and absorption cross sections, averaged over the energies of all of the neutrons, with the total number of neutrons emitted per fission, to create a single figure of merit which can be used for comparing different assemblies (see for example "One-Speed Transport Theory" in *Nuclear Reactor Theory* by George I. Bell and Samuel Glasstone, Van Nostrand Reinhold, 1970).

Fission neutrons start life with a Maxwell-Boltzmann statistical energy distribution, the same distribution law as gas molecules in a room, although here the temperature is close to 1 MeV (nearly 10 billion degrees C). The distribution looks like this:

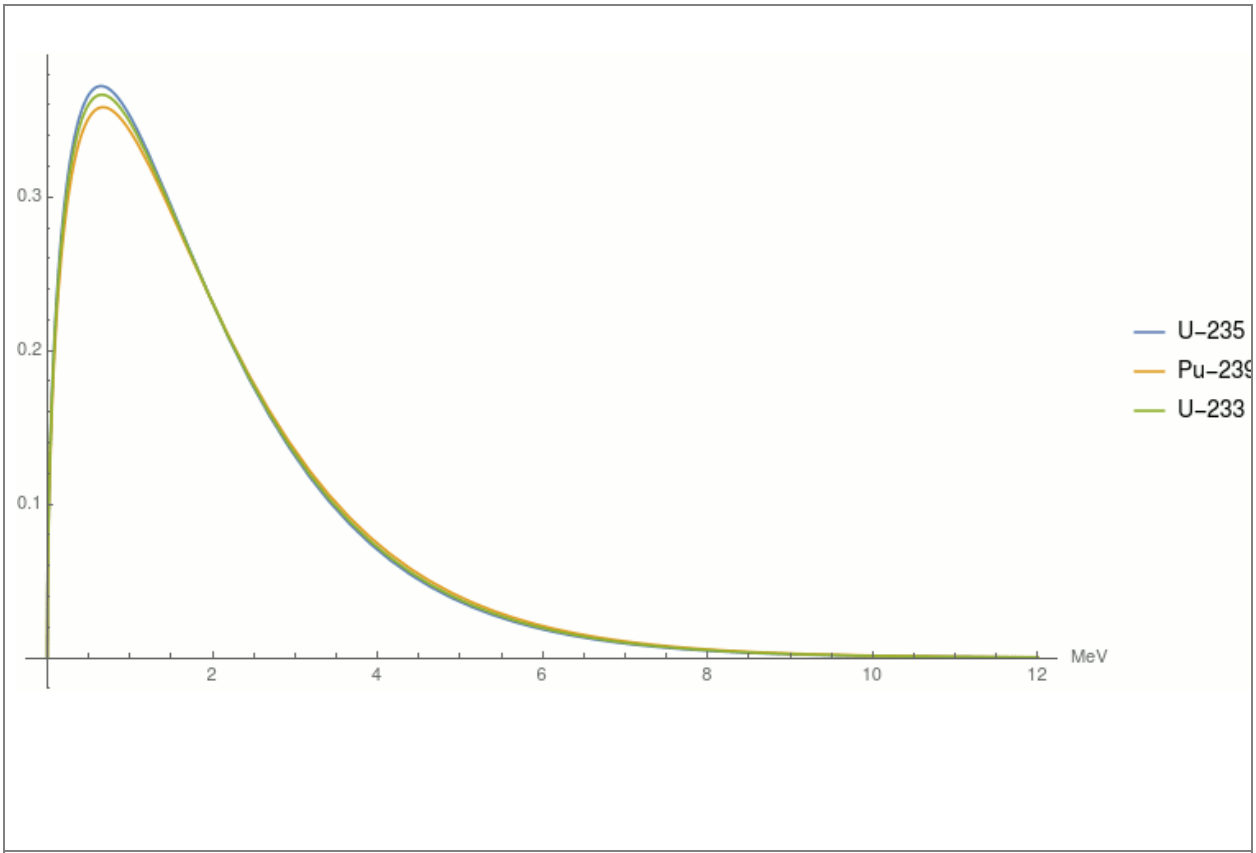

Probability density functions of the fission neutron energies of U-235, Pu-239 and U-233.

The most likely fission neutron energy is around 0.7 MeV, but the average value, due to the long tail to the right, is very close to 2 MeV.

Nuclear cross sections measure the likelihood of a nuclear reaction occurring, and are the effective geometric area the nucleus presents for that reaction. Larger cross sections are reactions that are more likely. Cross sections are measured in the non-SI unit called a "Barn", which is equal to 10^-24 cm^2.

The basic idea is this, when a neutron interacts with an atom we can think of it as consisting of two steps:

1. the neutron is "absorbed" by the collision; and
2. zero or more neutrons are emitted.

If the interaction is ordinary neutron capture, then no neutron is emitted from the collision. If the interaction is a scattering event, then one neutron is emitted. If the interaction is a fission event, then the average number of neutrons produced per fission is emitted (this average number is often designated by nu). By combining these we get the average number of neutrons produced per collision (also called the number of secondaries), designated by c:

Eq. 4.1.1-1

$$c = (cross_scatter + cross_fission*avg_n_per_fission - cross_absorption)/cross_total$$

the total cross section, cross_total, is equal to:

Eq. 4.1.1-2

$$cross_total = cross_scatter + cross_fission + cross_absorb$$

There are two types of scattering - elastic and inelastic.

Elastic scattering is the same process seen in the macroscopic world with billiard balls. Inelastic scattering only happens at high energies and in this case the model of the neutron being absorbed by the collision is more than conceptual, it really is absorbed froming a highly excited compound nucleus which then emits a neutron, and one of more gamma rays in the process. Elastic scattering causes energy loss, which is how moderators in reactors work, but neutrons bouncing off of heavy nuclei lose very little energy. For Pu-239 an elastic scattering event reduces the energy by only 0.84% on average.

Inelastic scattering, which converts a large part of the neutron kinetic energy into gamma rays, causes large energy losses, reducing the average fission spectrum neutron energy from 2 MeV to 1 MeV after one average inelastic scattering event. With inelastic scattering the conceptual model of scattering consisting of an absorption followed by emission is literally true. An unstable *compound nucleus* is formed in a very high energy state, but instead of undergoing fission it reemits a neutron at a lower energy, and disposes of the excess by emitting a gamma ray.The effect of inelastic scattering in plutonium is depicted below with the continuous distribution discretized into 50 logarithmically equal energy bands. The blue curve is the same distribution as shown above, but looks a bit different since it is plotted on a log-linear scale (the energy scale is logarithmic, the vertical probability axis is linear).

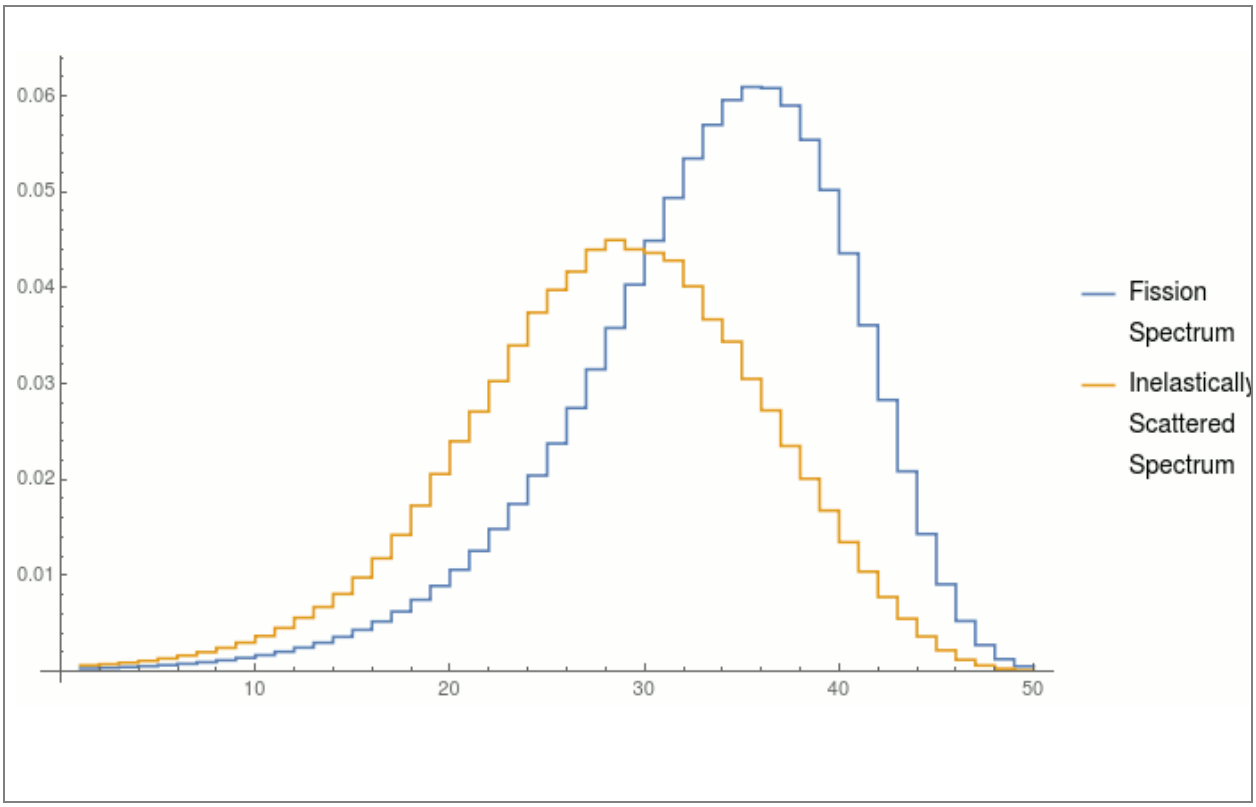

This loss of energy is important since it slows the neutron down, and fission explosions are very strongly affected by the speed of the reaction. The energy released varies by the third power of the reaction rate in fact. The cross sections also depend on the energy of the neutron, but this is a smaller effect than the energy velocity change.

Inelastic scattering causes another deviation from assumptions of diffusion theory, which supposes that all scattering events are isotropic. Inelastic scattering is not isotropic, but has a directional distribution peaked in the forward direction of neutron travel. But since all the other neutron "producing" events are isotropic, the inelastic scattering peak will also tend to be oriented isotropically. It is an established fact the diffusion theory, despite many simplifications of the real physical process, gives a very good general description of the behavior of fission weapons.

The total neutron mean free path (MFP), the average distance a neutron will travel before undergoing a collision, is given by:

Eq. 4.1.1-3

MFP = 1/(cross_total * N)

where N is the number of atoms per unit volume, determined by the density.

In computing the effective reactivity of a system we must also take into account the rate at which neutrons are lost by escape from the system. This rate is measured by the number of neutrons lost per collision. For a given geometry, the rate is determined by the size of the system in MFPs. Put another way, for a given geometry and degree of reactivity, the size of the system as measured in MFPs, is determined only by the parameter c. The higher the value of c, the smaller the assembly will be.

An indication of the effect of c on the size of a critical assembly can be gained by the following table of critical radii (in MFPs) for bare (unreflected) spheres:

Table 4.1.1-1. Critical Radius (r_c) vs Number of Secondaries (c)

Value of c	**r_c (critical radius in MFP)**
1.00	infinite
1.02	12.027
1.05	7.277
1.10	4.873
1.20	3.172
1.40	1.985
1.45	1.827
1.50	1.694
1.60	1.476

Many treatments of fission weapon physics ignore scattering and discuss only the fission mean free path and the fission neutron multiplicity (number of neutrons emitted per fission). This is fine when only considering the time-dependent behavior of the system, but scattering is crucial for understanding the critical size of the system, or the effect of a reflector. A nuclide with exactly the same fission cross section and multiplicity will have a smaller critical mass if its scattering cross section is higher. One can envision this by imagining a nuclide with an extremely high scattering cross section, in that case it is very hard for a neutron to escape as it scatters many times, and has a long "random walk" and so the critical mass must be very small.

If the composition, geometry, and reactivity of a system are specified then the size of a system in MFPs is fixed. From Eq. 4.1.1-3 we can see that the physical size or scale of the system (measured in centimeters, say) is inversely proportional to its density. Since the mass of the system is equal to volume*density, and volume varies with the cube of the radius, we can immediately derive the following scaling law:

Eq. 4.1.1-4

$$mcrit_c = mcrit_0/(rho/rho_0)^2 = mcrit_0/C^2$$

That is, the critical mass of a system is inversely proportional to the square of the density. C is the degree of compression (density ratio). This scaling law applies to bare cores, it also applies to cores with a surrounding reflector, if the reflector density has an identical degree of compression. This is usually not the case in real weapon designs, a higher degree of compression generally being achieved in the core than in the reflector.

An approximate relationship for this is:

Eq. 4.1.1-5

$$mcrit_c = mcrit_0/(C_c^1.2 * C_r^0.8)$$

where C_c is the compression of the core, and C_r is the compression of the reflector. Note that when C_c = C_r, then this is identical to Eq. 4.1.1-4. For most implosion weapon designs (since C_c > C_r) we can use the approximate relationship:

Eq. 4.1.1-6

$$mcrit_c = mcrit_0/C_c^1.7$$

These same considerations are also valid for any other specified degree of reactivity, not just critical cores.

Fission explosives depend on a very rapid release of energy. We are thus very interested in measuring the rate of the fission reaction. This is done using a quantity called the effective multiplication rate or "alpha". The neutron population at time t is given by:

Eq. 4.1.1-7

$$N_t = N_0 e^{(\alpha t)}$$

Alpha thus has units of 1/t, and the neutron population will increase by a factor of e (2.71...) in a time interval equal to 1/alpha. This interval is known as the "time constant" (or "e-folding time") of the system, t_c. The more familiar concept of "doubling time" is related to alpha and the time constant simply by:

Eq. 4.1.1-8

$$\text{doubling_time} = (\ln 2)/\alpha = (\ln 2) t_c$$

Alpha is often more convenient than t_c or doubling times since its value is bounded and continuous: zero at criticality; positive for supercritical systems; and negative for subcritical systems. The time constant goes to infinity at criticality. The term "time constant" seems unsatisfactory for this discussion though since it is hardly constant, t_c continually changes during reactivity insertion and disassembly. Therefore I will henceforth refer to the quantity 1/alpha as the "multiplication interval".

Alpha is determined by the reactivity (c and the probability of escape), and the length of time it takes an average neutron (for a suitably defined average) to traverse an MFP. If we assume no losses from the system then a maximum possible value for alpha can be calculated by:

Eq. 4.1.1-9

$$\alpha_{max} = (1/\tau)(c - 1) = (v_n/\text{total_MFP})(c - 1)$$

where tau is the average neutron lifetime between collisions; and v_n is the average neutron velocity. The average velocity for neutrons directly emitted by fission is 1.8x10^9 cm/sec (corresponding to the average energy of 2 MeV), but due to inelastic scattering an average velocity of 1.4x10^9 cm/sec (1 MeV) is a better number (this is what J. Carson Mark uses in his paper *Explosive Properties of Reactor-Grade Plutonium*). The "no losses" assumption is an idealization. It provides an upper bound for reaction rates, and provides a good indication of the relative reaction rates in different materials. For extremely supercritical assemblies, consisting of many critical masses, neutron losses could actually become negligible and approach the alphas

given below, but these conditions are not generally found in nuclear weapons, which do not usually require such degrees of supercriticality.

The factor c - 1 used above is the "neutron number", it represents the average neutron excess per collision. In real systems there is always some leakage, when this leakage is taken in account we get the "effective neutron number" which is always less than c - 1. When the effective neutron number is zero the system is exactly critical.

4.1.2 Nuclear Properties of Fissile Materials

The "temperature" of the fission neutron distribution varies very little among the various fissile nuclides, as can be seen the chart above.

Below is a table of the cross sections for the three fissile nuclides that are of primary interest, using three different values for the neutron energy to give a sense of how much variation there is in the energy range of interest. Also included is a calculation of the average cross sections for a material of considerable practical interest, weapon grade plutonium (WG-Pu) of a composition typical for the last phase of U.S. nuclear weapons production (0.012% Pu-238, 93.8% Pu-239, 5.8% Pu-240, 0.35% Pu-241, 0.022% Pu-242), alloyed with gallium (1.0% by weight, 3.5% molar), chosen since it represents the most common material found in modern nuclear weapons. The cross section for this material is given for 1 MeV, the best value to characterize fission weapon cores.

Table 4.1.2-1 Reaction Cross Sections (Barns, 10^-24 cm^2)

Material:	**U-233**			**U-235**			**Pu-239**			**WG-Pu Alloy**
Reaction Type	**1 MeV**	**2 MeV**	**Fiss. Spectrum**	**1 MeV**	**2 MeV**	**Fiss. Spectrum**	**1 MeV**	**2 MeV**	**Fiss. Spectrum**	**1 MeV**
Total	6.780	7.171	7.651	6.800	7.148	7.628	6.970	7.226	7.835	6.856
Elastic	3.860	3.688	4.474	3.650	3.452	4.451	3.665	3.390	4.396	3.663
Inelastic	1.002	1.484	1.208	1.843	2.349	1.858	1.545	1.853	1.581	1.544

Total Scattering	4.862	5.172	5.682	5.493	5.802	6.309	5.210	5.243	5.977	5.207
Fission	1.839	1.959	1.910	1.199	1.287	1.920	1.734	1.975	1.911	1.673
Absorption	0.079	3.981	0.059	0.108	0.059	0.087	0.026	0.008	0.053	0.0275

Since we are reducing all of these features of the system down to a single parameter (c) both elastic and inelastic scattering are treated the same, absorption also just reduces the effective number of secondaries.

A few observations about this table. One might suppose that the fission spectrum averaged cross section values, which use the actual initial neutron energy distribution, or 2 MeV which is the average value of this distribution, would be the best cross section values to use. But you will notice that the fission spectrum cross section of inelastic scattering and fission are roughly similar, which means that before they terminate their path in fission a large share neutrons will have scattered inelastically about once (this is almost exactly true for U-235). So an average neutron spends about half its flight at its original energy, and half after losing energy in inelastic scattering. The average value of neutron energy in the system then turns out to be about 1 MeV, and this cross section is the best of the three for our purpose here.

The actual value of alpha at a given density is the result of many interacting factors: the relative neutron density and cross sections values as a function of neutron energy, weighted by neutron velocity which in turn is determined by the fission neutron energy spectrum modified by the effects of both moderation and (principally) inelastic scattering.

Ideally the value of alpha should be determined by "integral experiments", that is, measured directly in the fissile material where all of these effects will occur naturally. Calculating tau and alpha from differential cross section measurements, adjusted neutron spectrums, etc. is fraught with potential error (although the data now available to the weapons states minimizes this problem).

In the table below are values of c, total mean free path length, and the maximum alpha at maximum uncompressed densities at 1 MeV. Compression to above normal density (achievable factors range up to 3 in weapons) reduces the MFPs, and increase alphas (and the *relative*, or scaled, physical dimensions of the system) proportionately. Note especially that the WG-Pu gallium alloy used in weapons under normal conditions is the lower density delta phase of plutonium, but collapses into the high density alpha phase under the low pressure (as high explosive detonation pressures go) of only 20 kilobars (2 GPa)

Table 4.1.2-2 Fissile Material Properties (1 MeV)						
Isotope	**c**	**Cross_total (barns)**	**Total MFP (cm)**	**Density**	**Alpha (1/microsecond)**	**t_double (nanosecond)**
U-233	1.41	6.780	3.08	18.93	190	3.69
U-235	1.25	6.800	3.07	18.93	112	6.05
Pu-239	1.50	6.970	2.87	19.86	252	2.86
WG-Pu Alloy	1.46	6.856	3.620	15.60	178	3.90

Values of c and total MFP can be easily calculated for mixtures of materials as I have done for WG-Pu. In real fission weapons (unboosted) effective values for alpha are typically in the range 25-250 (doubling times of 2.8 to 28 nanoseconds) due to the effects of neutron leakage in real systems.

Notice that using Table 4.1.1-1 and Table 4.1.2-2 we can immediately estimate the critical mass of a bare plutonium sphere:

mass_crit = [(2*1.694*2.867 cm)^3]*(Pi/6)*19.86 g/cm^3 = 9530 grams The published figure for the pure isotope is about 10.2 kg, if this is thought of in terms of critical radii the difference is only 7%. Since significant amounts of physics are left out of this treatment, this is not a bad agreement.

All nations interested in nuclear weapons technology have performed integral experiments to measure alpha, but published data is sparse and in general is limited to the immediate region of criticality. Collecting data for systems at high densities requires extremely difficult high explosive experiments, and data for high alpha systems can only be done in actual nuclear weapon tests.

Some integral alpha data is available for systems near prompt critical. The most convenient measurements are of the *negative* alpha value for fast neutron chain reactions at delayed criticality. Since at prompt critical alpha is exactly zero, the ratio of the magnitude of this delayed critical measurement to the fraction of fission neutrons that are delayed allows the alpha value to be calculated. These were the only sort of alpha measurements available to the Manhattan Project for the design of the first atomic bombs.

The most informative values are from the Godiva and Jezebel unreflected reactor experiments. These two systems used bare metal weapon grade cores, so the properties of weapons material was being measured directly. Godiva consisted of oralloy (93.71 wt% U-235, 5.24 wt% U-238, 1.05 wt% U-234), Jezebel of weapon-grade delta-phase plutonium alloy (94.134 wt% Pu-239, 4.848 wt% Pu-240, 1.018 wt% gallium):

Table 4.1.2-2 Properties of Bare Critical Metal Assemblies

Mass, Density, and Measured Alpha are at Delayed Critical (D.C.)

Assembly Name	Material	Critical Mass kg	Density (1/microsec)	Meas. Alpha Fraction	Del. Neutron (1/microsec)	Calc. "Alpha"
Godiva	Oralloy	52.25	18.71	-1.35	0.0068	199
Jezebel	WG-Pu	16.45	15.818	-0.66	0.0023	287

In a weapon the lifetime of a neutron, the length of time it exists in the system until it causes fission, or escapes, is the reciprocal of the weapon alpha. Most of the time the neutron will cause fission. In the bare critical assemblies referenced above we know, since they are just critical that for every fission all but one neutron will escape. Thus for plutonium, with a neutron multiplicity of 3.01, the escaping fraction is 2.01/3.01 = 67%, and for U-235 with a neutron multiplicity of 2.52, the escaping fraction is 1.52/2.52 = 60%. This means that the calculated alpha for the system is dominated by the escaping neutrons. Extracting the value for the underlying maximum alpha value requires a significant degree of modeling and calculation to remove the escaping fraction from the calculation.

The distribution of the neutron flux in Jezebel can be calculated through neutron diffusion and is show below.

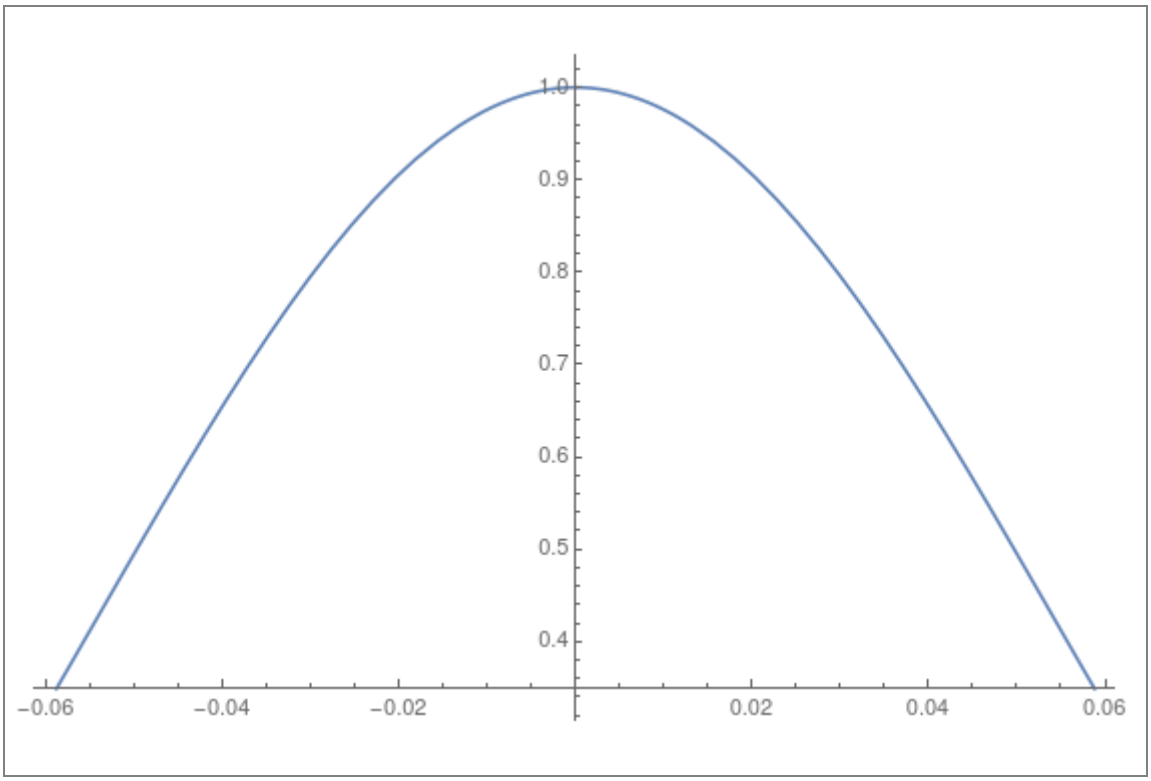

The effective value of alpha (the actual multiplication rate), taking into account neutron leakage, varies with the size of the system. If the system radius R = r_c, then it is exactly one critical mass (m = M_crit), and alpha is zero. The more critical masses present, the closer alpha comes to the limiting value. This can be estimated from the relation:

Eq. 4.1.2-1

alpha_eff = alpha_max*[1 - (r_c/R)^2]

= alpha_max*[1 - (M_crit/m)^(2/3)]

This approximates the change in neutron loss from the system, which declines as the radius increases beyond the critical radius, as more mass is added, i.e. at constant density. If the increase in the r_c/R ratio is due to compression factor C then alpha_max also increases to alpha_max*C.

4.1.3 Distribution of Neutron Flux and Energy in the Core

Since neutron leakage occurs at the surface of a critical or supercritical core, the strength of the neutron flux is not constant throughout the core. Since the rate of energy release at any point in the core is proportional to the flux at that point, this also affects the energy density throughout the core. This is a matter of some significance, since it influences weapon efficiency and the course of events in terminating the divergent fission chain reaction.

4.1.3.1 Flux Distribution in the Core

For a bare (unreflected) critical spherical system, the flux distribution is given by:

Eq. 4.1.3.1-1

flux(r) = max_flux * Sin(Pi*r/(r + 0.71*MFP))/(Pi*r/(r + 0.71*MFP))

(using the diffusion approximation) where Sin takes radians as an argument.

If we measure r in MFPs, then by referring to Table 4.1.1-1 we can relate the flux distribution to the parameter c. Computing the ratio between the flux at the surface of the critical system, and the maximum flux (in the center) we find:

Table 4.1.3-1 Relative Flux at Surface

Value of c	**r_c (critical radius in MFP)**
1.00	0.0 (at the limit)
1.02	0.0587
1.05	0.0963
1.10	0.1419
1.20	0.2117
1.40	0.3182
1.60	0.4018

This shows that as c increases, the flux distribution becomes flatter with less drop in the flux near the surface.

The flux distribution function above applies only to bare critical systems. If the system is supercritical, then the flux distribution becomes flatter, since neutron production over-balances loss. The greater the value of alpha for the system, the flatter it becomes. The addition of a neutron reflector also flattens the distribution, even for the same degree of reactivity (the radius is smaller in this case). The flux distribution function is useful though, since the maximum rate of fission occurs at the moment when the core passes through second criticality (on the way to disassembling, see below).

4.1.3.2 Energy Distribution in the Core

As long as the geometry doesn't change, the relative flux distribution remains the same throughout the fission process. The fission reaction rate at any point in the core is proportional to the flux. The net burnup of fissile material (and total energy release) is determined by the reaction rate integrated over time.

This indicates that the degree of burnup (the efficiency of utilization) varies throughout the core. The outer layers of material will be fissioned less efficiently than the material near the center. The steeper the drop off in flux the greater this effect will be. We can thus expect less efficient utilization of fissile material in small cores, and in materials with low values of c. From the relatively low value of c for U-235 compared to U-233 and Pu-239, we can expect that U-235 will be used less efficiently. This is observed in pure fission tests, the difference being about 15% in nominal yield (20 kt) pure fission designs.

The energy density (energy content per unit volume) in any region of the core is determined not only by the total energy produced in that region, but also by the flow of heat in to and out from the region.

The energy present in the core rises by a factor of e (2.71...) every multiplication interval (neglecting any losses from the surface). Nearly all of the energy present has thus been produced in the last one or two multiplication intervals, which in a high alpha system is a very short period of time (10 nanoseconds or less). There is not much time for heat flow to significantly alter this energy distribution.

Close to the end point of the fission process, the energy density in the core is so high that significant flow can occur. Since most of the energy is present as a photon gas the dominant mechanism is radiation (photon) heat transport, although electron kinetic heat transport may be significant as well. This heat flow can be modelled by the diffusion approximation just like neutron transport, but in this case estimating the photon mean free path (the opacity of the material) is quite difficult. A rough magnitude estimate for the photon MFP is a few millimeters.

The major of effect of energy flow is the loss of energy from a layer about 1 photon mean free path thick (referred to as one optical thickness) at the surface of the core. In a bare core this cooling can be quite dramatic, but the presence of a high-Z tamper (which absorbs and re-emits energy) greatly reduces this cooling. Losses also occur deeper in the core, but below a few

photon MFPs it becomes negligible. Otherwise, there is a significant shift in energy out of the center of the core that tends to flatten the energy distribution.

The energy density determines the temperature and pressure in the core, so there is also a variation in these parameters. Since the temperature in radiation dominated matter varies with the fourth power of the energy density, the temperature distribution is rather flat (except near the surface perhaps). The pressure is proportional to the energy density, so it varies in similar degree.

4.1.4 History of a Fission Explosion

To clarify the issues governing fission weapon design it is very helpful to understand the sequence of events that occurs in every fission explosion. The final event in the process - disassembly - is especially important since it terminates the fission energy release and thus determines the efficiency of the bomb.

4.1.4.1 Sequence of Events

Several distinct physical states can be identified during the detonation of a fission bomb. In each of these states a different set of physical processes dominates.

4.1.4.1.1 Initial State

Before the process that leads to a fission explosion is initiated, the fissile material is in a subcritical configuration. Reactivity insertion begins by increasing the average density of the configuration in some way.

4.1.4.1.2 Delayed Criticality

When the density has increased just to the point that a neutron population in the mass is self-sustaining, the state of delayed criticality has been achieved. Although nearly all neutrons produced by fission are emitted as soon as the atom splits (within 10^-14 sec or so), a very small proportion of neutrons (0.65% for U-235, 0.25% for Pu-239) are emitted by fission fragments with delays of up to a few minutes. In delayed criticality these neutrons are required to maintain the chain reaction. These long delays mean that power level changes can only occur slowly. All nuclear reactors operate in a state of delayed criticality. Due to the slowness of neutron multiplication in this state it is of no significance in nuclear explosions, although it is important for weapon safety considerations.

4.1.4.1.3 Prompt Criticality

When reactivity increases to the point that prompt neutrons alone are sufficient to maintain the chain reaction then the state of prompt criticality has been reached. Rapid multiplication can occur after this point. In bomb design the term "criticality" usually is intended to mean "prompt criticality". For our purposes we can take the value of alpha as being zero at this point. The reactivity change required to move from delayed to prompt criticality is quite small (for plutonium

the prompt and delayed critical mass difference is only 0.80%, for U-235 it is 2.4%), so in practice the distinction is unimportant. Passage through prompt criticality into the supercritical state is also termed "first criticality".

4.1.4.1.4 Supercritical Reactivity Insertion

The insertion time of a supercritical system is measured from the point of prompt criticality, when the divergent chain reaction begins. During this phase the reactivity climbs, along with the value of alpha, as the density of the core continues to increase. Any insertion system will have some maximum degree of reactivity which marks the end of the insertion phase. This phase may be terminated by reaching a plateau value, by passing the point of maximum reactivity and beginning to spontaneously deinsert, or by undergoing explosive disassembly.

4.1.4.1.5 Exponential Multiplication

This phase may overlap supercritical insertion to any degree. Any neutrons introduced into the core after prompt criticality will initiate a rapid divergent chain reaction that increases in power exponentially with time, the rate being determined by alpha. If exponential multiplication begins before maximum reactivity, and insertion is sufficiently fast, there may be significant increases in alpha during the course of the chain reaction. Throughout the exponential multiplication phase the cumulative energy released remains too small to disrupt the supercritical geometry on the time scale of the reaction. Exponential multiplication is always terminated by explosive disassembly. The elapsed time from neutron injection in the supercritical state to the beginning of explosive disassembly is called the "incubation time".

4.1.4.1.6 Explosive Disassembly

The bomb core is disassembled by a combination of internal expansion that accelerates all portions of the core outward, and the "blow-off" or escape of material from the surface, which generates a rarefaction wave propagating inward from the surface. The drop in density throughout the core, and the more rapid loss of material at the surface, cause the neutron leakage in the core to increase and the effective value of alpha to decline.

The speed of both the internal expansion and surface escape processes is proportional to the local speed of sound in the core. Thus disassembly occurs when the time it takes sound to traverse a significant fraction of the core radius becomes comparable to the time constant of the chain reaction. Since the speed of sound is determined by the energy density in the core, there is a direct relationship between the value of alpha at the time of disassembly and the amount of energy released. The faster is the chain reaction, the more efficient is the explosion.

As long as the value of alpha is positive (the core is supercritical) the fission rate continues to increase. Thus the peak power (energy production rate) occurs at the point where the core drops back to criticality (this point is called "second criticality"). Although this terminates the divergent chain reaction, and exponential increase in energy output, this does not mean that significant power output has ended. A convergent chain reaction continues the release of

energy at a significant, though rapidly declining, rate for a short time afterward. 30% or more of the total energy release typically occurs after the core has become sub-critical.

4.1.4.2 The Disassembly Process

The internal expansion of the core is caused by the existence of an internal pressure gradient. The escape of material from the surface is caused by an abrupt drop in pressure near the surface, allowing material to expand outward very rapidly. Both of these features are present in every fission bomb, but the degree to which each contributes to disassembly varies.

Consider a spherical core with internal pressure declining from the center towards the surface. At any radius r within the core the pressure gradient is dP/dR. Now consider a shell of material centered at r, that is sufficiently thin so that the slope of the pressure gradient does not change appreciably across it. The mass of the shell is determined by its area, density, and thickness:

$$m = \text{thickness} * \text{area} * \text{density}$$

The outward force exerted on the shell is determined by the pressure difference across the shell and the shell area:

$$F = dP/dR * \text{thickness} * \text{area}$$

From Newton's second law of motion we know that acceleration is related to force and mass by:

$$a = F/m$$

so:

$$a = (dP/dR * \text{thickness} * \text{area})/(\text{thickness} * \text{area} * \text{density})$$

$$= (dP/dR)/\text{density}$$

If density is constant in the core, then the outward acceleration at any point is proportional to the pressure gradient; the steeper the gradient, the greater the acceleration. The kinetic energy acquired comes at the expense of the internal energy of the expanding material.

The limiting case of a steep pressure gradient is a sudden drop to zero. In this case the acceleration is infinite, the internal energy of the material is completely converted to kinetic energy instantaneously and it expands outwards at constant velocity (escape velocity). The edge of the pressure drop propagates back into the material as a rarefaction wave at the local speed of sound. The pressure at the leading edge of the expanding material (moving in the

opposite direction at escape velocity) is zero. The pressure discontinuity thus immediately changes into a continuous pressure change of steadily diminishing slope. See Section 3.6.1.1 Release Waves for more discussion of this process.

In a bare core, thermal radiation from the surface causes a large energy loss in a surface layer about one optical thickness deep. Since energy lost from the core by thermal radiation cannot contribute to expansion, this has the effect of delaying disassembly. It does create a very steep pressure gradient in the layer however, and a correspondingly high outward acceleration. Deeper in the core, the pressure gradient is much flatter and the acceleration is lower. After the surface layer has expanded outward by a few times its original thickness, it has acquired considerable velocity, and the surface pressure drop rarefaction has propagated a significant distance back into the core. At this point the pressure and density profile of the core closely resembles the early stages of expansion from an instantaneous pressure drop, the development of the profile having been delayed slightly by the time it took the surface to accelerate to near escape velocity.

A bomb core will typically be surrounded by a high-Z tamper. A layer of tamper (about one optical thickness deep) absorbs the thermal radiation emitted by the core and is heated by it. As its temperature increases, this layer begins to radiate energy back to the core, reducing the core's energy loss. In addition, the heating also generates considerable pressure in the tamper layer. The combined effect of reduced core surface cooling, and this external pressure is to create a much more gradual pressure drop in the outer layer of the core and a correspondingly reduced acceleration.

The expanding core and heated tamper layer creates a shock wave in the rest of the tamper. This has important consequences for the disassembly process. The rarefaction wave velocity is not affected by the presence of the tamper, but the rate at which the density drops after arrival of the rarefaction wave is strongly affected. The rate of density drop is determined by the limiting outward expansion velocity, this is in turn determined by the shock velocity in the tamper. The denser the tamper the slower the shock, and the slower the density decrease behind the rarefaction wave. In any case the shock velocity in the tamper is much slower than the escape velocity of expansion into a vacuum. The disassembly of a tamped core thus more closely resembles one dominated by internal expansion rather than surface escape.

4.1.4.3 Post Disassembly Expansion

The expanding core creates a radiation dominated shock wave in the tamper that compresses it by at least a factor of 7, and perhaps as high as 16 due to ionization effects. This pileup of high density material at the shock front is called the "snow plow" effect. By the time this shock has moved a few centimeters into the tamper, the rarefaction wave will have reached the center of the core and the entire core will be expanding outward uniformly.

The basic structure of the early fireball has now developed, consisting of a thin highly compressed shell just behind the shock front containing nearly all of the mass that has been shocked and heated so far. This shell travels outward at nearly the same velocity as the shock

front. The volume inside this shell is a region of very low density. Temperature and pressure behind the shock front is essentially uniform though since nearly all of the energy present is contained in the radiation field (i.e. it exists as a photon gas). Since the shock wave is radiation dominated, the front does not contain an abrupt pressure jump. Instead there is a transition zone with a thickness about equal to the radiation mean free path in the high-Z tamper material (typically a few millimeters). In this zone the temperature and pressure climb steadily to their final value.

This overall explosion structure remains the same as the shock expands outward until it reaches a layer of low-Z material (a beryllium reflector, or the high explosive).

The transition zone marking the shock front remains thin as long as the shock is travelling through opaque high-Z material. Low-Z material becomes completely ionized as it is heated, and once it is completely ionized it is nearly transparent to radiation and is no longer efficiently heated. When the shock front emerges at the boundary of the high-Z tamper and the low-Z material, it spits into two regions. A radiation driven shock front moves quickly away from the high-Z surface, bleaching the low-Z material to transparency. This faster shock front only creates a partial transition to the final temperature and pressure. The transition is completed by a second shock, this one a classical mechanical shock, driven by the opaque material.

4.1.5 Fission Weapon Efficiency

Fundamental to analyzing the design of fission bombs is understanding the factors that influence the efficiency of the explosion - the percentage of fissile material actually fissioned. The efficiency and the amount of fissile material present determine the amount of energy released by the explosion - the bomb's yield.

I have organized my discussion of design principles around the issue of efficiency since it is the most important design characteristic of any fission device. Any weapon designer must have a firm grasp on the expected efficiency in order to make successful yield predictions, and a firm grasp on the factors affecting efficiency is required to make design tradeoffs.

In the discussion below (and in later subsections as well) I assume that the system under discussion is spherically symmetric, and of homogenous density, unless otherwise stated. Spherical symmetry is the simplest geometry to analyze, and also happens to be the preferred geometry for efficient nuclear weapons.

4.1.5.1 Efficiency Equations

It is intrinsically difficult to accurately predict the performance of a particular design from fundamental physical principles alone. To make good predictions on this basis requires sophisticated computer simulations that include hydrodynamic, radiation, and neutronic effects. Even here it is very valuable to have actual test data to use for calibrating these simulation models.

Nuclear weapon programs have historically relied heavily on extrapolating tested baseline designs using scaling laws like the efficiency equations I discuss below, especially in the early years of development. These equations are derived from idealized models of bomb core behavior and consequently have serious limitations in making absolute efficiency estimates. The predictions of the Theoretical Section at Los Alamos underestimated the yield of the first atomic bomb by a factor of three; an attempts a few years later to recompute the bomb efficiency using the best models, physical data, and computers available at the time led to a yield overestimate by a factor of two.

From the description of core disassembly given above we can see that two possible idealizations are possible for deriving convenient efficiency equations:

- uniform expansion of a core; and
- surface escape from a core initially at constant pressure.

The basic approach is to model how quickly the core expands to the point of second criticality. To within a constant scaling factor, this fixes the efficiency of the explosion.

In the first modelling approach, the state of second criticality is based on the average density of the entire core. In the second approach, second criticality is based on the surface loss of excess critical masses from a residual core which remains at constant initial density.

The first efficiency equation to be developed was the Bethe-Feynmann equation, prepared by Hans Bethe and Richard Feynman at Berkeley in 1942 based on the uniform expansion model. A somewhat different efficiency equation was presented by Robert Serber in early 1943 at Los Alamos, which was also based on uniform expansion but also explicitly included the exponential growth in energy release (which the Bethe-Feynmann equation did not). A problem with these derivations is that to keep the resultant formulas relatively simple, they assume that the expanding core remains at essentially constant density during deinsertion, which is only true (even approximately) when the degree of supercriticality is small.

For the purposes of this FAQ I have taken the second approach for deriving an efficiency equation, using the surface escape model. This model has the advantage that the residual core remains at constant density regardless of the degree of supercriticality. Comparing it to the other efficiency equations provides some insight into the sensitivity of the assumptions in the various models.

4.1.5.1.1 The Serber Efficiency Equation Revisited

Let us first consider the factors that affect the efficiency of a homogenous untamped supercritical mass. In this system, disassembly begins as fissile material expands off the core's surface into a vacuum. We make the following simplifying assumptions:

- Reactivity deinsertion is complete when the rarefaction wave reaches the critical radius of the core;

- The value of alpha does not change until the rarefaction wave reaches the critical radius, then it goes to zero;
- The temperature is uniform through the core, and no energy is lost.

If r is the initial outer radius, and r_c is the critical radius, then the reaction halts when:

Eq. 4.1.5.1.1-1

$$\text{Integral}[c_s(t)\, dt] = r - r_c$$

where c_s(t) is the speed of sound at time t.

If kinetic pressure is negligible compared to radiation pressure (this is true in all but extremely low yield explosions), then:

Eq. 4.1.5.1.1-2

$$c_s(t) = [(E(t)*\text{gamma})/(3*V*\text{rho})]^{0.5}$$

where E(t) is the cumulative energy produced by the reaction, V is the volume of the core, and rho is its density.

We also have:

Eq. 4.1.5.1.1-3

$$E(t) = (E1/(c - 1)) * e^{(\text{alpha}*t)}$$

where E1 is a constant that gives the energy yield per fission (E1 = 2.88 x 10^-4 erg/fission). Thus:

Eq. 4.1.5.1.1-4

$$\text{Eff}(t) = E(t)/E_{\text{total}} = (E1/((c - 1)*E_{\text{total}})) * e^{(\text{alpha}*t)}$$

where Eff(t) is the efficiency at time t, and E_total is the energy yield at 100% efficiency.

Thus:

Eq. 4.1.5.1.1-5

r - r_c = Integral[(E(t)*gamma/(3*V*rho))^0.5 dt]

= (gamma*E1/(3*M*(c-1)))^0.5 * Integral[e^(alpha*t/2)dt]

= (gamma*E1/(3*M*(c-1)))^0.5 * 2/alpha * e^(alpha*t/2)

where M is the fissile mass.

Rearranging and squaring we get:

Eq. 4.1.5.1.1-6

e^(alpha*t) = (r - r_c)^2 * ((3M*(c-1))/(gamma*E1)) * (alpha^2)/4

Substituting into the efficiency equation:

Eq. 4.1.5.1.1-7

Eff(t) = [3*alpha^2 * M * (r - r_c)^2]/(4*gamma*E_total)

If E2 is a constant equal to fission energy/gram in ergs (7.25 x 10^17 erg/g for Pu-239), and gamma is equal to 4/3 for a photon gas, then:

Eq. 4.1.5.1.1-8

Eff(t) = [9*alpha^2 * (r - r_c)^2]/(16*E2)

We can observe at this point that efficiency is determined by the actual value of alpha and the difference between the actual radius of the assembly, and the radius of the mass just sufficient to keep the chain reaction going. Note that it is the values of these parameters WHEN DISASSEMBLY ACTUALLY OCCURS that are relevant.

Now using r = r_c(1 + delta) so that (r - r_c) = delta*r_c, we get:

Eq. 4.1.5.1.1-9

Eff(t) = [9*alpha^2 * delta^2 * r_c^2]/(16*E2)

If we let tau = (total_MFP/v_n) then:

Eq. 4.1.5.1.1-10

alpha_max = (v_n/total_MFP)*(c - 1) = (c - 1)/tau

and

Eq. 4.1.5.1.1-11

alpha_eff = ((c - 1)/tau)*[1 - (1/(1 + delta)^2)]

Now:

Eq. 4.1.5.1.1-12

Eff(t) = ((c-1)/tau)^2 * 9/(16*E2) * r_c^2 * delta^2 *[1-(1/(1+ delta)^2)]^2

= ((c-1)/tau)^2 * 9/(16*E2) * r_c^2 *[delta - (delta/(1+ delta)^2)]^2

In the range of 0 < delta < 1 (up to 8 critical masses), the expression

[delta - (delta/(1+ delta)^2)]^2

is very close to 0.6*delta^3, giving us:

Eq. 4.1.5.1.1-13

Eff(t) = 0.338*((c-1)/tau)^2 * r_c^2/E2 * delta^3

= 0.338/E2 * alpha_max^2 * r_c^2 * delta^3

This last equation is identical with the equation derived by Robert Serber in the spring of 1943 and published in *The Los Alamos Primer*, except that his constant is 0.667 (i.e. gives efficiencies 1.98 times higher). Serber derived his efficiency equation from rough dynamical considerations without using a hydrodynamic model of disassembly and admits that his result is 2-4 time higher than the true value. This is consistent with the above derivation.

Both the equation given above and Serber's equation differ significantly from the Bethe-Feynman equation however, which gives an efficiency relationship of:

Eq. 4.1.5.1.1-14

Eff = (1/(gamma - 1)E2) * alpha_max^2 * r_c^2 *

(delta*(1 + 3*delta/2)^2)/(1 + delta)

after reformulating to equivalent terms. This is a much more linear relationship between delta and efficiency, than the cubic relationship of Serber. Due to the crudeness of all of these derivations, the significance of this difference cannot be assessed at present.

Equation 4.1.5.1.1-13 shows that efficiency is proportional to the square of the maximum multiplication rate of the material, and the critical radius (also due to material properties), and is the cube of the excess critical radius excess delta.

Extending to larger values, we can approximate it in the range 1 < delta < 3 (up to 64 critical masses), with the expression:

Eq. 4.1.5.1.1-15

Eff(t) = 0.338/E2 * alpha_max^2 * r_c^2 * delta^(7/3)

4.1.5.1.2 The Density Dependent Efficiency Equation

The efficiency equations given above leave something to be desired for evaluating fission weapon designs. I have included it to assist in making comparisons with the available literature, but I will give it a different form below.

The choice of fissile materials available to a weapon designer is quite limited, and the nuclear and physical properties of these materials are fixed. It is desirable then to separate these factors from the factors that a designer can influence - namely, the mass of material present, and the density achieved. The density is of particular interest since it is the only factor that changes in a given design during insertion. Understanding how efficiency changes with density is essential to understanding the problem of predetonation for example.

Returning to equation Eq. 4.1.5.1.1-8:

Eff(t) = [9*alpha^2 * (r - r_c)^2]/(16*E2)

we want to reformulate it so that it consists of two parts, one that does not depend on density, and one that depends only on density.

Let the composition and mass of the system be fixed. We will normalize the radius and density so that they are expressed relative to the system's critical state. If rho_crit and r_crit are the values for density and radius of the critical state, and rho_rel and r_rel are the values of the system that we want to evaluate:

Eq. 4.1.5.1.2-1

$$\rho_{rel} = \rho_{actual}/\rho_{crit}$$

and

Eq. 4.1.5.1.2-2

$$r_{rel} = r_{actual}/r_{crit}$$

When the system is exactly critical, rho_rel = 1 and r_rel = 1. Of course we are interested in states where rho_rel > 1, and r_rel < 1. We can relate r_rel to rho_rel:

Eq. 4.1.5.1.2-3

$$r_{rel} = (1/\rho_{rel})^{1/3} * r_{crit}$$

Using this notation, and letting alpha_max_c be the value of alpha_max at the critical state density, we can write:

$$\alpha = \alpha_{max_c} * \rho_{rel} * (1 - (r_c/r_{rel})^2)$$

In this case r_c refers to the effective critical radius at density rho_rel not rho_crit; that is, r_c IS NOT r_crit. Instead it is equal to r_crit/rho_rel. Using this, and the relation for r_rel above, we can eliminate r_crit:

Eq. 4.1.5.1.2-4

$$\alpha = \alpha_{max_c} * \rho_{rel} * (1 - ((1/\rho_{rel})/(1/\rho_{rel})^{1/3}))^2$$

$$= \alpha_{max_c} * \rho_{rel} * (1 - (\rho_{rel})^{-4/3})$$

Substituting into the efficiency equation:

Eq. 4.1.5.1.2-5

Eff = (9/16*E2) * alpha^2 * (r_rel - r_c)^2

we get:

Eq. 4.1.5.1.2-6

Eff = (9/(16*E2))*(alpha_max_c*rho_rel*(1 - (rho_rel)^(-4/3)))^2 *

(r_rel - r_c)^2

Splitting constant and density dependent factors between two lines:

Eq. 4.1.5.1.2-7

Eff = (9/(16*E2)) * alpha_max_c^2 *

rho_rel^2 * (1-(rho_rel)^(-4/3))^2 * (r_rel - r_c)^2

We can eliminate r_rel and r_c, replacing them with expressions of rho_rel and r_crit:

Eq. 4.1.5.1.2-8

r_rel - r_c = (1/rho_rel)^(1/3) * r_crit) - (r_crit/rho_rel)

= ((1/rho_rel)^(1/3) - (1/rho_rel)) * r_crit

Substituting again:

Eq. 4.1.5.1.2-9

Eff = (9/(16*E2)) * alpha_max_c^2 * r_crit^2 *

rho_rel^2 * (1-(rho_rel)^(-4/3))^2 * ((1/rho_rel)^(1/3)-(1/rho_rel))^2

Recall that the rho_rel, the relative density, is not generally the compression ratio compared to normal density. This is true only if amount of fissile material in the system is exactly one critical mass at normal density (as was approximately true in the Fat Man bomb). For "sub-crit" systems, rho_rel is smaller than the actual compression of the material since compressive work

is required to raise the initial sub-critical system to the critical state. For a system consisting of more than one critical mass (at normal density), rho_rel is higher than the actual compression.

By looking in turn at each of the density dependent terms we can gain insight into the significance of the efficiency equation. First note that alpha_max_c is a fundamental property of the fissile material and does not change, even though it is system dependent (being normalized to the critical density of the system).

The term (rho_rel^2) is introduced by the reduction of the MFP with increasing density and contributes to enhanced efficiency at all values of rho_rel.

The term (1-(rho_rel)^(-4/3)))^2 represents the effect of neutron leakage. At rho_rel=1 the value is 0. It has a limiting value of 1 when rho_rel is high, i.e. no leakage occurs. As this term approaches one, and leakage becomes insignificant, it ceases to be a significant contributor to further efficiency enhancement.

The term ((1/rho_rel)^(1/3)-(1/rho_rel))^2 describes the distance the rarefaction wave must travel to shut down the reaction. At rho_rel=1 it is 0. It initially increases rapidly, but soon slows down at reaches a maximum at about rho_rel = 5.196. Thereafter it declines slowly. This signifies that fact that once the critical radius of the system at rho_rel is small compared to the physical radius no further efficiency gain is obtained from this source. Instead further increases in density simply reduce the scale of the system, allowing faster disassembly.

We can provide some approximations for the efficiency equation to make the overall effect of density more apparent.

In the range of 1 < rho_rel < 2 it is approximately:

Eq. 4.1.5.1.2-10

$$Eff = (9/(16*E2)) * alpha_max_c^2 * r_crit^2 * ((rho_rel - 1)^3)/8$$

In the range of 2 < rho_rel < 4.5 it is approximately:

Eq. 4.1.5.1.2-11

$$Eff = ((9/(16*E2)) * alpha_max_c^2 * r_crit^2 * ((rho_rel - 1)^(2.333))/8$$

In the range of 4 < rho_rel < 8 it is approximately:

Eq. 4.1.5.1.2-12

$$Eff = (9/(16*E2)) * alpha_max_c^2 * r_crit^2 * ((rho_rel - 1)^(1.8))/5$$

4.1.5.1.3 The Mass and Density Dependent Efficiency Equation

The maximum degree of compression above normal density that is achievable is limited by technology. It is of interest then to consider how the amount of material present affects efficiency at a given level of compression, since it is the other major parameter that a designer can manipulate.

To examine this we would like to reintroduce an explicit term for mass. To do this we renormalize the equation to a fixed standard density rho_0 (the uncompressed density of the fissile material), and use rho_0 and the corresponding value of the critical mass M_c to replace the scale parameter r_crit. Thus:

Eqs. 4.1.5.1.3-1 through 4.1.5.1.3-5

alpha_max_crit = alpha_max_0 * (rho_crit/rho_0)

m_rel = m/M_c

rho_crit = rho_0/m_rel^(1/2)

rho_rel = rho/rho_crit = (rho/rho_0)*m_rel^(1/2)

r_crit = ((m/rho_crit)*(3/(2Pi)))^(1/3)

= (m*m_rel^(1/2)/rho_0)^(1/3) * (3/2Pi)^(1/3)

= (m^(3/2)/(M_c^(1/2) * rho_0))^(1/3) * (3/2Pi)^(1/3)

= m^(1/2) * (M_c^(1/2) * rho_0)^(-1/3) * (3/2Pi)^(1/3)

Assuming the density rho >= rho_crit, we get:

Eq. 4.1.5.1.3-6

Eff = (9/(16*E2))*(3/(2Pi))^(2/3) * alpha_max_0^2 *

(rho_crit/rho_0)^2 * (rho/rho_crit)^2 *

(m^(1/2) * (M_c^(1/2) * rho_0)^(-1/3))^2 *

(1-((rho_0/rho)^(4/3) * m_rel^(-2/3)))^2 *

(((rho_0/rho)^(1/3) * m_rel^(-1/6)) - ((rho_0/rho) * m_rel^(-1/2)))^2

Simplifying:

Eq. 4.1.5.1.3-7

Eff = (9/(16*E2))*(3/(2Pi))^(2/3) * alpha_max_0^2 *

(rho/rho_0)^2 * m/(M_c^(1/3) * rho_0^(2/3)) *

(1-((rho_0/rho)^(4/3) * m_rel^(-2/3)))^2 *

m_rel^(-1) * (((rho_0 * m_rel)/rho)^(1/3) - (rho_0/rho))^2

Then:

Eq. 4.1.5.1.3-8

Eff = (9/(16*E2))*(3/(2Pi))^(2/3) * alpha_max_0^2 * m/(M_c^(1/3)) * (M_c/m)

(rho^2)/(rho_0^(8/3)) * (1 - ((rho_0/rho)^(4/3) * m_rel^(-2/3)))^2 *

(((rho_0 * m_rel)/rho)^(1/3) - (rho_0/rho))^2

And finally:

Eq. 4.1.5.1.3-9

Eff = (9/(16*E2))*(3/(2Pi))^(2/3) * alpha_max_0^2 * M_c^(2/3) *

(rho/(rho_0^(4/3)))^2 * (1 - ((rho_0/rho)^(4/3) * m_rel^(-2/3)))^2 *

(((rho_0 * m_rel)/rho)^(1/3) - (rho_0/rho))^2

The first line of this equation consists entirely of constants, some of them fixed by the choice of material and reference density. From the next two lines it is clear that the density dependency is the same. The effect of increasing the mass of the system is to modestly reduce leakage and retard disassembly.

4.1.5.1.4 The Mass Dependent Efficiency Equation

It is useful to also have an equation that considers only the effect of mass. Including this as the only variable allows presenting a simplified form that makes the effect of varying the mass in a particular design easier to visualize. Also in gun-type designs no compression occurs, so the chief method of manipulating yield is by varying the mass of fissile material present.

Taking the mass and density dependent equation, we can set the density to a fixed nominal value, rho, and then simplify. Let rho = rho_0:

Eq. 4.1.5.1.4-1

Eff = (9/16*E2)*(3/2Pi)^(2/3) * alpha_max_0^2 * M_c^(2/3) *

(rho_0/(rho_0^(4/3))^2 *(1 - ((rho_0/rho_0)^(4/3) * m_rel^(-2/3)))^2 *

(((rho_0 * m_rel)/rho_0)^(1/3) - (rho_0/rho_0))^2

= (9/16*E2)*(3/2Pi)^(2/3) * alpha_max_0^2 * M_c^(2/3) *

rho_0^(-2/3) * (1 - m_rel^(-2/3))^2 * ((m_rel)^(1/3) - 1)^2

Since M_c/rho_0 is the volume of a critical assembly (m_rel = 1):

Eq. 4.1.5.1.4-2

Eff = (9/16*E2)*(3/2Pi)^(2/3) * alpha_max_0^2 * vol_crit^(2/3) *

(1 - m_rel^(-2/3))^2 * ((m_rel)^(1/3) - 1)^2

And finally:

Eq. 4.1.5.1.4-3

Eff = (9/16*E2)*(2^(2/3)) * alpha_max_0^2 * r_crit^2 *

(1 - m_rel^(-2/3))^2 * ((m_rel)^(1/3) - 1)^2

Again the top line consists of numeric and material constants, the second of mass dependent terms. This equation shows that efficiency is zero when m_rel = 1, as expected. Efficiency is negligible when m_rel < 1.05, similar to the power of conventional explosives. It climbs very quickly however, increasing by a factor of 400 or so between 1.05 and 1.5, where efficiency becomes significant. The Little Boy bomb had m_rel = 2.4. If its fissile content had been

increased by a mere 16%, its yield would have increased by 75% (whether this could be done while maintaining a safe criticality margin is a different matter).

4.1.5.1.5 Limitations of the Efficiency Equations

These formulas provide good scaling laws, and a rough means to calculate efficiency. But we should return to the simplifying assumptions made earlier to understand their limitations.

It is obvious that alpha is not constant during disassembly. As material blows off, the size of the core and the value of alpha both decrease, which has a negative effect on efficiency. This is the most important factor not accounted for, and results in a lower effective coefficient in the efficiency equation.

The assumption about uniform temperature, and no energy loss is also not really true. The energy production rate in any region of the core is proportional to the neutron flux density. This density is highest in the center and lowest at the surface (although not dramatically so). Furthermore, the high radiation energy density in the core corresponds to a high radiation loss rate from the surface. Based on the Stefan-Boltzmann law it would seem that the loss rate from a bare core could eventually match the energy production rate. This doesn't really occur because of the high opacity of ionized high-Z material; thermal energy from inside the core cannot readily reach the surface. But by the same token, the surface can cool dramatically. Since core expansion starts at the surface, and the rate is determined by temperature, this surface cooling can significantly retard disassembly.

When scaling from known designs, most of these issues have little significance since the deviations from the theoretical model used for the derivations affects both system similarly.

The efficiency equations also breaks down at very small yields. To eliminate gamma from the equations I assumed that the core was radiation dominated at the time of disassembly. When yields drop to the low hundreds of tons and below, the value of gamma approximates that of a perfect gas which changes only the constant term in the equations, reducing efficiency by 20%. When yields drop to the ton range then the properties of condensed matter (like physical strength, heat of vaporization, etc.) become apparent. This tends to increase the energy release since these properties resist the expansion effects.

There is another factor that imposes an effective upper limit on efficiency regardless of other attempts to enhance yield. This is the decrease in fissile content of the core. The alert reader may have noticed that it is possible to calculate efficiencies that are greater than 1 using the equations. This is because energy release is represented as an exponentially increasing function of time without regard for the amount of energy actually present in the fissile material. At some point, the fact that the fission process depletes the fissile material present must have an effect on the progress of the chain reaction.

The limiting factor here is due to the dilution of the fissile material by the fission products. Most isotopes have roughly the same absorption cross section for fast neutrons, a few barns. The core initially consists of fissile material, but as the chain reaction proceeds each fission event

replaces one fissile nucleus with two fission product nuclei. When 50% of the material has fissioned, for every 100 initial fissile atoms there are now 50 remaining, and 100 non-fissile atoms, i.e. the fissile content has declined to only 33%. This parasitic absorption will eventually extinguish the reaction entirely, regardless of what yield enhancement techniques are used (generally at an efficiency substantially below 50%).

4.1.5.2 Effect of Tampers and Reflectors on Efficiency

So far I have been explicitly assuming a bare fissile mass for efficiency estimation. Of course, most designs surround the core with layers of material intended to scatter escaping neutrons back into the fissile mass, or to retard the hydrodynamic expansion.

I use the term "reflector" to refer to the neutron scattering properties of the surrounding material, and "tamper" to refer to the effect on hydrodynamic expansion. The distinction is logical because the two effects are fundamentally unrelated, and because the term tamper was borrowed from explosive blasting technique where it refers only to the containment of the blast. This distinction is not usually made in US weapons programs, from Manhattan Project on. The custom is to use "tamper" to refer to both effects, although "neutronic tamper" and "reflector" are used if the neutron reflection effect alone is intended.

4.1.5.2.1 Tampers

In the bare core, the fissile material that has been reached by the inward moving rarefaction wave expands outward very rapidly. In radiation dominated matter, expansion into a vacuum reaches a limiting speed of six times the local speed of sound in the material (this is the velocity at the outer surface of the expanding sphere of material). The density of matter behind the rarefaction front (which moves toward the center of the core) thus drops very rapidly and is almost immediately lost to the fission reaction.

If a layer of dense material surrounds the core then something very different occurs. The fissile material is not expanding into a vacuum, instead it has to compress and accelerate matter ahead of it. That is, it creates a shock wave. The expansion velocity of the core is then limited to the velocity of accelerated material behind the expanding shock front, which is close to the shock velocity itself. If the tamper and fissile core have similar densities, then this expansion velocity is similar to the speed of sound in the core and only 1/6 as fast as the unimpeded expansion velocity.

This confining effect means that the drop in alpha as disassembly proceeds is not nearly as abrupt as in a vacuum. It thus reduces the importance of the inaccurate assumption of constant alpha used in deriving the efficiency equation.

Another important effect is caused by the radiation cooling of the core. In a vacuum this energy is lost to free space. An opaque tamper absorbs this energy, and a layer of material one mean free path thick is heated to nearly the temperature and pressure of the core. The expansion shock wave then arises not at the surface of the core, but some distance away in the tamper (on the order of a few millimeters). A rarefaction wave must then propagate back to the surface of

the core before its expansion even begins. In effect, this increases size of the expansion distance term ((1/rho_rel)^(1/3)-(1/rho_rel))^2 in the efficiency equation.

4.1.5.2.2 Reflectors

In a bare core, any neutron that reaches the surface of the core is lost forever to the reaction. A reflector scatters the neutrons, a process that causes some fraction of them to eventually reenter the fissile mass (usually after being scattered several times). Its effect on efficiency then can be described simply by reducing the neutron leakage term (rho_rel)^(-4/3) by a constant factor, or by reducing the reference density critical mass terms.

The leakage or critical mass adjustments must take into account time absorption effects. This means that leakage cannot simply be reduced by the probability of a lost neutron eventually returning, and the reflected critical mass cannot be based simply on the steady state criticality value. For example when an efficiently reflected assembly is only slightly supercritical, then multiplication is dependent mostly (or entirely) on the reflected neutrons that reenter the core. On average each of these neutrons spends quite a lot of time outside the core before being scattered back in. The relevant value for alpha_max in this system is not the value for the fissile material, but is instead:

alpha _max = 1/(average neutron life outside of core)

This is likely to be at least an order of magnitude larger than the core material alpha_max value.

4.1.5.3 Predetonation

An optimally efficient fission explosion requires that the explosive disassembly of the core occur when the neutron multiplication rate (designated alpha) is at a maximum. Ideally the bomb will be designed to compress the core to this state (or close to it) before injecting neutrons to initiate the chain reaction. If neutrons enter the mass after criticality, but before this ideal time, the result is predetonation (or preinitiation): disassembly at a sub-optimal multiplication rate, producing a reduced yield.

How significant this problem is depends on the reactivity insertion rate. Something like 45 multiplication intervals must elapse before really significant amounts of energy are released. Prior to this point predetonation is not possible. The number of these intervals that occur during a period of time is obtained by integrating alpha over the period. When alpha is effectively constant it is simply alpha*t.

During insertion, alpha is not constant. When insertion begins its value is zero. If a neutron is injected early in insertion and insertion is slow, we can accumulate 45 multiplication intervals when alpha is still quite low. In this case a dramatic reduction in yield will occur. On the other hand, if it were possible for insertion to be so fast that full insertion is achieved before

accumulating enough multiplication intervals to disassemble the bomb then no predetonation problem would exist.

To evaluate this problem let us consider a critical system with initial radius r_0 undergoing uniform spherical compression, with the radius decreasing at a constant rate v, then alpha is:

Eq. 4.1.5.3-1

$$\alpha = \alpha_{max_0} * ((r_0/(r_0 - v*t))^3 - ((r_0 - v*t)/r_0))$$

Integrating, we obtain:

Eq. 4.1.5.3-2

$$Int[\alpha] = \alpha_{max_0}*(r_0^3/(2v*(r_0-v*t)^2) - (t-(v*t^2)/(2*rc)))$$

Which allows to compute the number of elapsed multiplication intervals between times t_1 and t_2.

For example, consider a system with the following parameters with a critical radius r = 4.5 cm, a radial implosion velocity v = 2.5x10^5 cm/sec, and alpha_max_0 = 2.8x10^8/sec. Figure 4.1.5.3-1 shows the accumulation of elapsed neutron multiplication intervals (Y axis) as implosion proceeds (seconds on X axis).

Figure 4.1.5.3-1. Elapsed Multiplication Intervals Vs Implosion Time

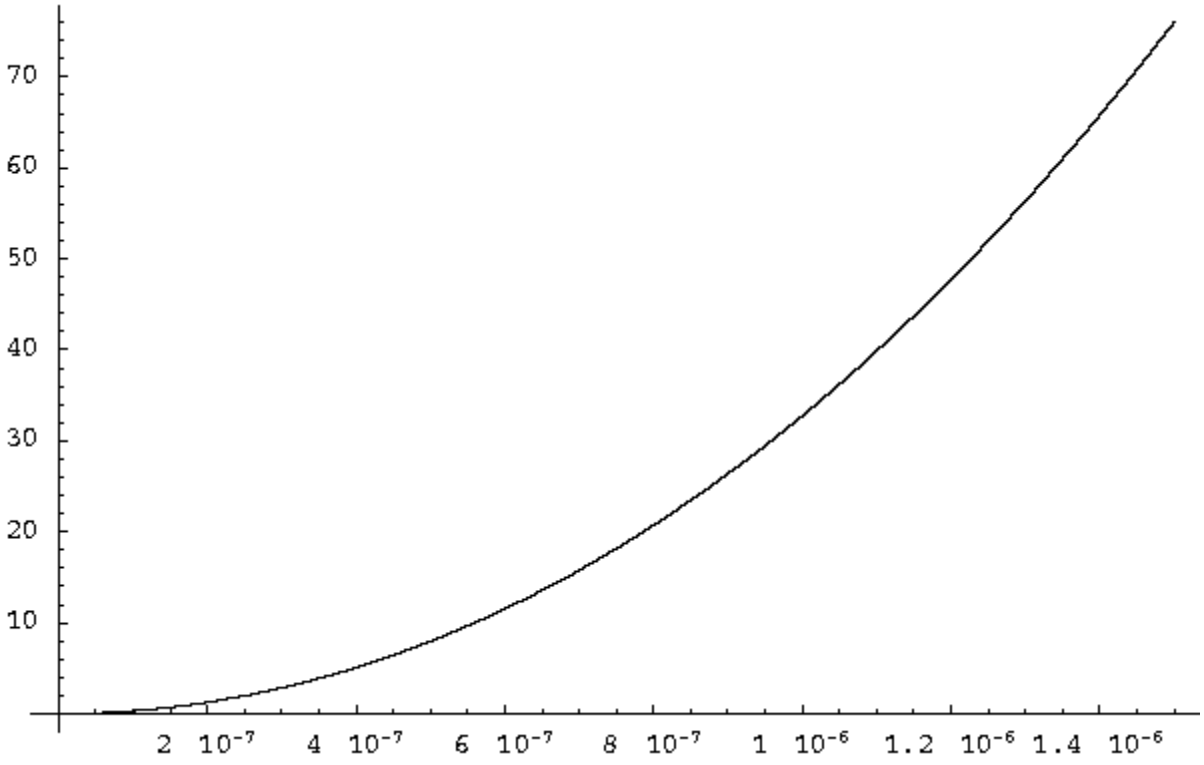

Recall that disassembly occurs when the speed of sound, c_s, integrated over the life of the chain reaction is equal to r - r_c, the difference between the outer radius and the critical radius. Since c_s is proportional to the square root of the energy released, it increases by a factor of e every 2 multiplication intervals. Disassembly thus occurs quite abruptly, effectively occurring over a period of two multiplication intervals. The condition for disassembly is thus:

Eq. 4.1.5.3-3

r(t) - r_c(t) = 2*c_s(t)/alpha(t) for some time t.

Since r - r_c is a polynomial function, and c_s is a transcendental (exponential) function, no closed form means of calculating t is possible. However these functions are monotonically increasing in the range of values of interest so numeric and graphical techniques can easily determine when the disassembly condition occurs. The value of alpha at that point then determines efficiency.

Taking our previous example (r = 4.5 cm, v = 2.5x10^5 cm/sec, alpha_max = 2.8x10^8/sec) we can plot the net implosion distance (r - r_c) and the integrated expansion distance (2*c_s/alpha) against the implosion time. This is shown in the log plot in Figure 4.1.5.3-2 for the period between 1 and 1.3 microseconds. Distance is in centimeters (Y axis) and time is in seconds (X axis). If a neutron is present at the beginning of insertion, we see that the disassembly condition

occurs at t = 1.25x10^-6 sec. At this point 52 multiplication intervals have elapsed, and the effective value of alpha is 8.6x10^7/sec. The corresponding yield is about 0.5 kt.

Figure 4.1.5.3-2. Implosion Distance and Expansion Distance Plotted Against Implosion Time

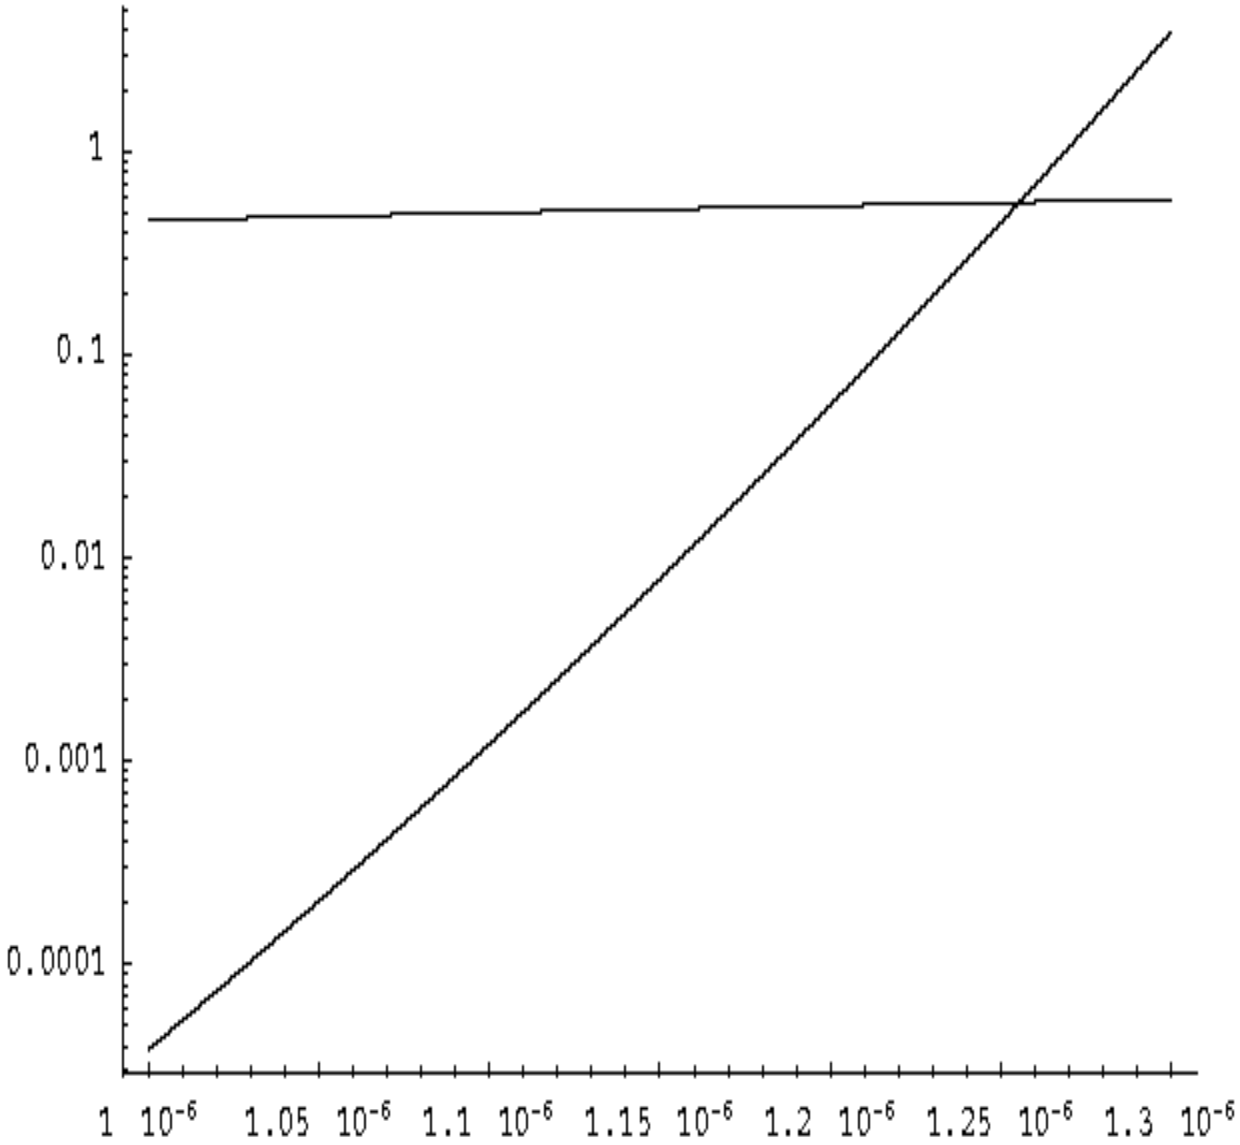

The parameters above approximately describe the Fat Man bomb. This shows that even in the worst case, neutrons being present at the moment of criticality, quite a substantial yield would have been created. Predetonation does not necessarily result in an insignificant fizzle. It is not feasible though to make a high explosive driven implosion system fast enough to completely defeat predetonation through insertion speed alone (radiation driven implosion and fusion boosting offer means of overcoming it however).

The likelihood of predetonation occurring depends on the neutron background, the average rate at which neutron injection events occur. I use the term "neutron injection event" instead of simply talking about neutrons for a specific reason: the major source of neutrons in a fission device is spontaneous fission of the fissile material itself (or of contaminating isotopes). Each spontaneous fission produces an average of 2-3 neutrons (depending on the isotope). However, these neutrons are all released at the same moment, and thus either a fission chain reaction is initiated at the moment, or they all very quickly disappear. Each fission is a single injection event, neutrons from other sources are uncorrelated and are thus individual injection events.

Now neutron injection during insertion is not guaranteed to initiate a divergent chain reaction. At criticality (alpha equals zero), each fission generates on average one fission in the next generation. Since each fission produces nu neutrons (nu is in the range of 2-3 neutrons, 2.9 for Pu-239), this means that each individual neutron has only 1/nu chance of causing a new fission. At positive values of alpha, the odds are better of course, but clearly we must consider then the probability that each injection actually succeeds in creating a divergent chain reaction. This probability is dependent on alpha, but since non-fission capture is a significant possibility in any fissile system, it does not truly converge to 1 regardless of how high alpha is (although with plutonium it comes close).

Near criticality the probability of starting a chain reaction (P_chain) for a single neutron is thus about 34% for plutonium, and 40% for U-235. Since spontaneous fission injects multiple neutrons, the P_chain for this injection event is high, about 70% for both Pu-239 and U-235.

If the average rate of neutron injection is R_inj, then the probability of initiating a chain reaction during an insertion time of length T is the Poisson function: Eqs. 4.1.5.3-4 P_init = 1 - e^((-T/R_inj)*P_chain) If T is much smaller than R_inj then this equation reduces approximately to P_init = (T/R_inj)*P_chain.

When T is much smaller than R_inj predetonation is unlikely, and the yield of the fission bomb (which will be the optimum yield) can be predicted with high confidence. As the ratio of T/R_inj becomes larger yield variability increases. When (T/R_inj)*P_chain is equal to ln 2 (0.693...) then the probability of predetonation and no predetonation is equal, although when predetonation occurs close to full assembly the yield reduction is small. As T/R_inj continues to increase predetonation becomes virtually certain. With a large enough value to T/R_inj the yield becomes predictable again, but this time it is the minimum yield that results when neutrons are present at the beginning of insertion. For an implosion bomb a typical spread between the optimum and minimum yields is something like 40:1.

In the Fat Man bomb the neutron source consisted of about 60 g of Pu-240, which produced an average of one fission every 37 microseconds. The probability of predetonation was 12% (from a declassified Oppenheimer memo), assuming an average P_chain of 0.7 we can estimate the insertion time at 6.7 microseconds, or 4.7 microseconds if P_chain was close to 1. The chance of large yield reduction was much smaller than this however. There was a 6% chance of a yield < 5 kt, and only a 2% chance of a yield < 1 kt. As we have seen, in no case would the yield have been smaller than 0.5 kt or so.

Spontaneous fission is not the only cause for concern, since neutrons can enter the weapon from outside. Natural neutron sources are not cause for concern, but in a combat situation very powerful sources of neutrons may be encountered - other nuclear weapons.

One kiloton of fission yield produces a truly astronomical number of excess neutrons - about 3x10^24, with a fluence of 1.5x10^10 neutrons/cm^2 500 m away. A kiloton of fusion yields 3-4 times as many. The fission reaction itself emits all of its neutrons in less than a microsecond, but due to moderation these neutrons arrive at distant locations over a much longer period of time. Most of them arrive in a pulse lasting a millisecond, but thermal neutrons can continue to arrive for much longer periods of time. This is not the whole problem though. Additional neutrons called "delayed neutrons" continue to be emitted for about a minute from the excited fission products. These amount to only 1% or so of the prompt neutrons, but this is still an average arrival rate of 2.5x10^6 neutrons/cm^2-sec for a kiloton of fission at 500 m. With weapons sensitive to predetonation, careful spacing of explosions in distance and time may be necessary. Neutron hardening - lining the bomb with moderating and neutron absorbing materials - may be necessary to hold predetonation problems to a tolerable level (it is virtually impossible to eliminate it entirely in this way).

4.1.6 Methods of Core Assembly

The principal problem in fission weapon design is how to rapidly assemble or compress the fissile material from a subcritical state to a supercritical one. Methods of doing this can be classified in two ways:

- Whether it is subsonic or supersonic; and
- Number of geometric axes along which compression occurs.

Subsonic assembly means that shock waves are not involved. Assembly is performed by adiabatic compression, or by continuous acceleration. As a practical matter, only one subsonic assembly scheme needs to be considered: gun assembly.

Supersonic assembly means that shock waves are involved. Shock waves cause instantaneous acceleration, and naturally arise whenever the very large forces required for extremely rapid assembly occur. The are thus the natural tools to use for assembly. Shocks are normally created by using high explosives, or by collisions between high velocity bodies (which have in turn been accelerated by high explosive shocks). The term "implosion" is generally synonymous with supersonic assembly. Most fission weapons have been designed with assembly schemes of this type.

Assembly may be performed by compressing the core along one, two, or three axes. One-D compression is used in guns, and plane shock wave compression schemes. Two and three-D compression are known as cylindrical implosion and spherical implosion respectively. Plane shock wave assembly might logically be called "linear Implosion", but this term has been usurped (in the US at any rate) by a variant on cylindrical implosion (see below). The basic

principles involved with these approaches are discussed in detail in Section 3.7, Principles of Implosion.

To the approaches just mentioned, we might add more some difficult to classify hybrid schemes such as: "pseudo-spherical implosion", where the mass is compressed into a roughly spherical form by convergent shock waves of more complex form; and "linear implosion" where a compressive shock wave travels along a cylindrical body (or other axially symmetric form - like an ellipsoid), successively squeezing it from one end to the other (or from both ends towards the middle). Schemes of this sort may be used where high efficiency is not called for, and difficult design constraints are involved, such as severe size or mass limitations. Hybrid combinations of gun and implosion are also possible - firing a bullet into an assembly that is also compressed.

The number of axes of assembly naturally affect the overall shape of the bomb. One-D assembly methods naturally tend to produce long, thin weapon designs; 2-D methods lead to disk-shaped or short cylindrical systems; and 3-D methods lead to spherical designs.

The subsections detailing assembly methods are divided in gun assembly (subsonic assembly) and implosion assembly (supersonic assembly). Even though it superficially resembles gun assembly, linear implosion is discussed in the implosion section since it actually has much more in common with other shock compression approaches.

The performance of an assembly method can be evaluated by two key metrics: the total insertion time and the degree of compression. Total insertion time (and the related insertion rate) is principally important for its role in minimizing the probability of predetonation. The degree of compression determines the efficiency of the bomb, the chief criteria of bomb performance. Short insertion times and high compression are usually associated since the large forces needed to produce one also tend to cause the other.

4.1.6.1 Gun Assembly

This was the first technique to be seriously proposed for creating fission explosions, and the first to be successfully developed. The first nuclear weapon to be used in war was the gun-type bomb called Little Boy, dropped on Hiroshima. Basic gun assembly is very simple in both concept and execution. The supercritical assembly is divided into two pieces, each of which is subcritical. One of these, the projectile, is propelled into the other, called the target, by the pressure of propellant combustion gases in a gun barrel. Since artillery technology is very well developed, there are really no significant technical problems involved with designing or manufacturing the assembly system.

The simple single-gun design (one target, one projectile) imposes limits on weapon, mass, efficiency and yield that can be substantially improved by using a "double-gun" design using two projectiles fired at each other. These two approaches are discussed in separate sections below. Even more sophisticated "complex" guns, that combine double guns with implosion are discussed in Hybrid Assembly techniques.

Gun designs may be used for several applications. They are very simple, and may be used when development resources are scarce or extremely reliability is called for. Gun designs are natural where weapons can be relatively long and heavy, but weapon diameter is severely limited - such as nuclear artillery shells (which are "gun type" weapons in two senses!) or earth penetrating "bunker busters" (here the characteristics of a gun tube - long, narrow, heavy, and strong - are ideal).

Single guns are used where designs are highly conservative (early US weapon, the South African fission weapon), or where the inherent penalties of the design are not a problem (bunker busters perhaps). Double guns are probably the most widely used gun approach (in atomic artillery shells for example).

4.1.6.1.1 Single Gun Systems

We might conclude that a practical limit for simple gun assembly (using a single gun) is a bit less than 2 critical masses, reasoning as follows: each piece must be less than 1 critical mass, if we have two pieces then after they are joined the sum must be less than 2 critical masses.

Actually we can do much better than this. If we hollow out a supercritical assembly by removing a chunk from the center like an apple core, we reduce its effective density. Since the critical mass of a system is inversely proportional to the square of the density, we have increased the critical mass remaining material (which we shall call the target) while simultaneously reducing its actual mass. The piece that was removed (which will be called the bullet) must still be a bit less than one critical mass since it is solid. Using this reasoning, letting the bullet have the limiting value of one full critical mass, and assuming the neutron savings from reflection is the same for both pieces (a poor assumption for which correction must be made) we have:

Eq. 4.1.6.1.1-1

$$M_c/((M - M_c)/M)^2 = M - M_c$$

where M is the total mass of the assembly, and M_c is the standard critical mass. The solution of this cubic equation is approximately M = 3.15 M_c. In other words, with simple gun assembly we can achieve an assembly of no more than 3.15 critical masses. Of course a practical system must include a safety factor, and reduce the ratio to a smaller value than this.

The weapon designer will undoubtedly surround the target assembly with a very good neutron reflector. The bullet will not be surrounded by this reflector until it is fired into the target, its effective critical mass limit is higher, allowing a larger final assembly than the 3.15 M_c calculated above.

Looking at U-235 critical mass tables for various candidate reflectors we can estimate the achievable critical mass ratios taking into account differential reflector efficiency. A steel gun barrel is actually a fairly good neutron reflector, but it will be thinner and less effective than the

target reflector. M_c for U-235 (93.5% enrichment) reflected by 10.16 cm of tungsten carbide (the reflector material used in Little Boy) is 16.5 kg, when reflected by 5.08 cm of iron it is 29.3 kg (the steel gun barrel of Little Boy was an average of 6 cm thick). This is a ratio of 1.78, and is probably close to the achievable limit (a beryllium reflector might push it to 2). Revising Eq. 4.1.6.1.1-1 we get:

Eq. 4.1.6.1.1-2

$$M_c/((M - (1.78\ M_c))/M)^2 = M - (1.78\ M_c)$$

which has a solution of M = 4.51 M_c. If a critical mass ratio of 2 is used for beryllium, then M = 4.88 M_c. This provides an upper bound on the performance of simple gun-type weapons.

Some additional improvement can be had by adding fast neutron absorbers to the system, either natural boron, or boron enriched in B-10. A boron-containing sabot (collar) around the bullet will suppress the effect of neutron reflection from the barrel, and a boron insert in the target will absorb neutrons internally thereby raising the critical mass. In this approach the system would be designed so that the sabot is stripped of the bullet as it enters the target, and the insert is driven out of the target by the bullet. This system was apparently used in the Little Boy weapon.

Using the M_c for 93.5% enriched U-235, the ratio M/M_c for Little Boy was (64 kg)/(16.5 kg) = 3.88, well within the limit of 4.51 (ignoring the hard-to-estimate effects of the boron abosrbers). It appears then that the Little Boy design (completed some six months before the required enriched uranium was available) was developed with the use of >90% enrichment uranium in mind. The actual fissile load used in the weapon was only 80% enriched however, with a corresponding WC reflected critical mass of 26.5 kg, providing an actual ratio of 64/26.5 = 2.4.

The mass-dependent efficiency equation shows that it is desirable to assembly as many critical masses as possible. Applying this equation to Little Boy (and ignoring the equation's limitations in the very low yield range) we can examine the effect of varying the amount of fissile material present:

1.05 80 kg

1.1 1.2 tons

1.2 17 tons

1.3 78 tons

1.4 220 tons

1.5 490 tons

1.6 930 tons

1.8 2.5 kt

2.0 5.2 kt

2.25 10.5 kt

2.40 15.0 kt LITTLE BOY

2.5 18.6 kt

2.75 29.6 kt

3.0 44 kt

3.1

If its fissile content had been increased by a mere 25%, its yield would have tripled.

The explosive efficiency of Little Boy was 0.23 kt/kg of fissile material (1.3%), compared to 2.8 kt/kg (16%) for Fat Man (both are adjusted to account for the yield contribution from tamper fast fission). Use of 93.5% U-235 would have at least doubled Little Boy yield and efficiency, but it would still have remained disappointing compared to the yields achievable using implosion and the same quantity of fissile material.

4.1.6.1.2 Double Gun Systems

Significant weight savings a possible by using a "double-gun" - firing two projectiles at each other to achieve the same insertion velocity. With all other factors being the same (gun length, projectile mass, materials, etc.) the mass of a gun varies with the fourth power of velocity (doubling velocity requires quadrupling pressure, quadrupling barrel thickness increases mass sixteen-fold). By using two projectiles the required velocity is cut by half, and so is the projectile mass (for each gun). On the other hand, to keep the same total gun length though, the projectile must be accelerated in half the distance, and of course there are now two guns. The net effect is to cut the required mass by a factor of eight. The mass of the breech block (which seals the end of the gun) reduces this weight saving somewhat, and of course there is the offsetting added complexity.

A double gun can improve on the achievable assembled mass size since the projectile mass is divided into two sub-critical pieces, each of which can be up to one critical mass in size. Modifying Eq. 4.1.6.1.1-1 we get:

Eq. 4.1.6.1.1-3

$$M_c/((M - 2M_c)/M)^2 = M - 2M_c$$

with a solution of $M = 4.88 M_c$.

Taking into account the effect of differential reflector efficiency we get mass ratios of ratios of 3.56 (tungsten carbide) and 4 (beryllium) which give assembled mass size limits of $M = 7.34 M_c$ and $M = 8 M_c$ respectively.

In 2016 Joseph Thompson suggested to me a more complex gun design that could increase the number of crits achievable to a very high level. If instead of a single solid piece being taken out of the supercritical assembly, the idea is that both the target and projectile consist of multiple concentric cylindrical shells that nest together to make a solid mass. Then on average each piece is 1/2 the density of the supercritical assembly, and thus 1/4 of the number of crits. Thus each piece separately can be slightly less than one crit. When a second piece is added to it, it doubles the mass, but also doubles the density, leading to a total of nearly *eight* crits.

An interesting aspect to his scheme is that since the two pieces are equivalent it makes it easy to reason about the insertion, or assembly, time problem - the fact that the mass becomes critical before the two pieces begin to insert or even meet. This is also addressed in "Section 4.1.6.1.3 Weapon Design and Insertion Speed" below.

Since the two pieces are of identical properties when they are adjacent (just before physical insertion begins) they are in effect a single half density piece of about two crits, but with a length of 2L, where L is the length of each piece. There is an adjustment, called the "shape factor", that must be made since this is not the optimal compact cylinder with an L/D ratio of 1, but a cylinder of L/D=2. Shape factor curves from criticality tests of highly reflected HEU show that the reduction here is 17%, so that we really have 1.70 crits.

For these two pieces not to form a critical assembly they must be separated. We can make an estimate of how large this separation must be by treating the separation as a reduction in density. For two critical masses to become one the density must decrease by a factor of 1/SQRT(1.72), or the the opposite ends must be 2*SQRT(1.70)L apart which means that the gap is 2*SQRT(1.70)L - 2L, or about 0.608L. Of course this increases the shape factor effect, but only by about 4.5, so the gap is really slightly less than this. Thus the entire insertion time during which predetonation could occur for this system is the time it takes to travel 1.6L.

This idea of pieces that are effectively homogenous low density nesting components that assemble like a puzzle to form a solid mass can be extended to a double gun and three pieces. While a scheme to support a set of two concentric cylindrical shells is easily imagined (supporting them on one end of the piece, how to do it with the central piece to allow insertion from both sides would be more of a trick. But assuming on has such a system, then each piece has 1/3 the mass, and 1/3 the density, so when the whole system is assembled you get to 27 crits! In this case the whole assembly will need a length of about 4.3L to avoid being critical, but the insertion gaps on either end are only modestly larger, about 0.65L.

Another variant of the double gun concept is to still only have two fissile masses - a hollow mass and a cylindrical core as in the single gun - but to drive them both together with propellant. One

possible design would be to use a constant diameter gun bore equal to the target diameter, with the smaller diameter core being mounted in a sabot. In this design the target mass would probably be heavier than the core/sabot system, so one end of the barrel might be reinforced to take higher pressures. Another more unusual approach would be to fire the target assembly down an annular (ring shaped) bore. This design appears to have been used in the U.S. W-33 atomic artillery shell, which is reported to have had an annular bore.

These larger assembled masses give significantly more efficient bombs, but also require large amounts of fissile material to achieve them. And since there is no compression of the fissile material, the large efficiency gains obtainable through implosive compression is lost. These shortcomings can be offset somewhat using fusion boosting, but gun designs are inherently less efficient than implosion designs when comparing equal fissile masses or yields.

4.1.6.1.3 Weapon Design and Insertion Speed

In addition to the efficiency and yield limitations, gun assembly has some other significant shortcomings:

First, guns tend to be long and heavy. There must be sufficient acceleration distance in the gun tube before the projectile begins insertion. Increasing the gas pressure in the gun can shorten this distance, but requires a heavier tube.

Second, gun assembly is slow. Since it desirable to keep the weight and length of the weapon down, practical insertion velocities are limited to velocities below 1000 m/sec (usually far below). The diameter of a core is on the order of 15 cm, so the insertion time must be at least a 150 microseconds or so.

In fact, achievable insertion times are much longer than this. Taking into account only the physical insertion of the projectile into the core underestimates the insertion problem. As previously indicated, to maximize efficiency both pieces of the core must be fairly close to criticality by themselves. This means that a critical configuration will be achieved before the projectile actually reaches the target. The greater the mass of fissile material in the weapon, the worse this problem becomes. With greater insertion distances, higher insertion velocities are required to hold the probability of predetonation to a specified value. This in turn requires greater accelerations or acceleration distances, further increasing the mass and length of the weapon.

In Little Boy a critical configuration was reached when the projectile and target were still 25 cm apart. The insertion velocity was 300 m/sec, giving an overall insertion time of 1.35 milliseconds.

Long insertion times like this place some serious constraints on the materials that can be used in the bomb since it is essential to keep neutron background levels very low. Plutonium is excluded entirely, only U-235 and U-233 may be used. Certain designs may be somewhat sensitive to the isotopic composition of the uranium also. High percentages of even-numbered isotopes may make the probability of predetonation unacceptably high.

The 64 kg of uranium in Little Boy had an isotopic purity of about 80% U-235. The 12.8 kg of U-238 and U-234 produced a neutron background of around 1 fission/14 milliseconds, giving Little Boy a predetonation probability of 8-9%. In contrast to the Fat Man bomb, predetonation in a Little Boy type bomb would result in a negligible yield in nearly every case.

The predetonation problem also prevents the use of a U-238 tamper/reflector around the core. A useful amount of U-238 (200 kg or so) would produce a fission background of 1 fission/0.9 milliseconds.

Gun-type weapons are obviously very sensitive to predetonation from other battlefield nuclear explosions. Without hardening, gun weapons cannot be used within a few of kilometers of a previous explosion for at least a minute or two.

Attempting to push close to the mass limit is risky also. The closer the two masses are to criticality, the smaller the margin of safety in the weapon, and the easier it is to cause accidental criticality. This can occur if a violent impact dislodges the projectile, allowing it to travel toward the target. It can also occur if water leaks into the weapon, acting as a moderator and rendering the system critical (in this case though a high yield explosion could not occur).

Due to the complicated geometry, calculating where criticality is achieved in the projectile's travel down the barrel is difficult, and calculating the effective value of alpha vs time as insertion continues is even harder. Computation intensive Monte Carlo techniques are required, but with general purpose Monte Carlo transport codes now available, and modern computers, this is not longer a serious problem. In the development of Little Boy these things had to be extrapolated from measurements made in scale models.

4.1.6.1.4 Initiation

Once insertion is completed, neutrons need to be introduced to begin the chain reaction. One route to doing this is to use a highly reliable "modulated" neutron initiator, an initiator that releases neutrons only when triggered. The sophisticated neutron pulse tubes used in modern weapons are one possibility. The Manhattan Project developed a simple beryllium/polonium 210 initiator named "Abner" that brought the two materials together when struck by the projectile.

If neutron injection is reliable, then the weapon designer does not need to worry about stopping the projectile. The entire nuclear reaction will be completed before the projectile travels a significant distance. On the other hand, if the projectile can be brought to rest in the target without recoiling back then an initiator is not even strictly necessary. Eventually the neutron background will start the reaction unaided.

A target designed to stop the projectile once insertion is complete is called a "blind target". The Little Boy bomb had a blind target design. The deformation expansion of the projectile when it impacted on the stop plate of the massive steel target holder guaranteed that it would lodge firmly in place. Other designs might add locking rings or other retention devices. Because of the use of a blind target design, Little Boy would have exploded successfully without the Abner initiators. Oppenheimer only decided to include the initiators in the bomb fairly late in the

preparation process. Even without Abner, the probability that Little Boy would have failed to explode within 200 milliseconds was only 0.15%; a delay as long as one second was vanishingly small - 10^-14.

Atomic artillery shells have tended to be gun-type systems, since it is relatively easy to make a small diameter, small volume package this way (at the expense of large amounts of U-235). Airbursts are the preferred mode of detonation for battlefield atomic weapons which, for an artillery shell travelling downward at several hundred meters per second, means that initiation must occur at a precise time. Gun-type atomic artillery shells always include polonium/beryllium initiators to ensure this.

4.1.6.2 Implosion Assembly

High explosive driven implosion assembly uses the ability of shock waves to instantaneously compress and accelerate material to high velocities. This allows compact designs to rapidly compress fissile material to densities much higher than normal on a time scale of microseconds, leading to efficient and powerful explosions. The speed of implosion is typically several hundred times faster than gun assembly (e.g. 2-3 microseconds vs. 1 millisecond). Densities twice the normal maximum value can be reached, and advanced designs may be able to do substantially better than this (compressions of three and four fold are often claimed in the unclassified literature, but these seem exaggerated). Weapon efficiency is typically an order of magnitude better than gun designs.

The design of an implosion bomb can be divided into two parts:

1. The shock wave generator: the high explosive system that generates an initial shock wave of the appropriate shape;
2. The implosion hardware: the system of inert materials that is driven by the shock wave, which consists of the nuclear explosive materials, plus any tampers, reflectors, pushers, etc. that may be included.

The high explosive system may be essentially unconfined (like that in the Fat Man bomb), but increased explosive efficiency can be obtained by placing a massive tamper around the explosive. The system then acts like a piston turned inside out, the explosive gases are trapped between the outer tamper and the inner implosion hardware, which is driven inward as the gases expand. The added mass of the tamper is no doubt greater than the explosive savings, but if the tamper is required anyway (for radiation confinement, say) then it adds to the compactness of the design.

If you have not consulted Section 3.7 Principles of Implosion, it may be a good idea to do so.

4.1.6.2.1 Energy Required for Compression As explained in Section 3.4 Hydrodynamics, shock compression dissipates energy in three ways:

1. through work done in compressing the shocked material,
2. by adding kinetic energy to the material (accelerating it), and

3. by increasing the entropy of the material (irreversible heating).

Only the first of these is ultimately desirable for implosion, although depending on the system design some or all of the kinetic energy may be reclaimable as compressive work. The energy expended in entropic heating is not only lost, but also makes the material more resistant to further compression.

Shock compression always dissipates some energy as heat, and is less efficient than gentle isentropic (constant entropy) compression. Examining the pressure and total energy required for isentropic compression thus provides a lower bound on the work required to reach a given density.

Below are curves for the energy required for isentropic and shock compression of uranium up to a compression factor of 3. For shock compression only the energy the appears as internal energy (compression and heating) are included, kinetic energy is ignored.

Figure 4.1.6.2.1-1. Required Energy for Shock and Isentropic Compression of Uranium

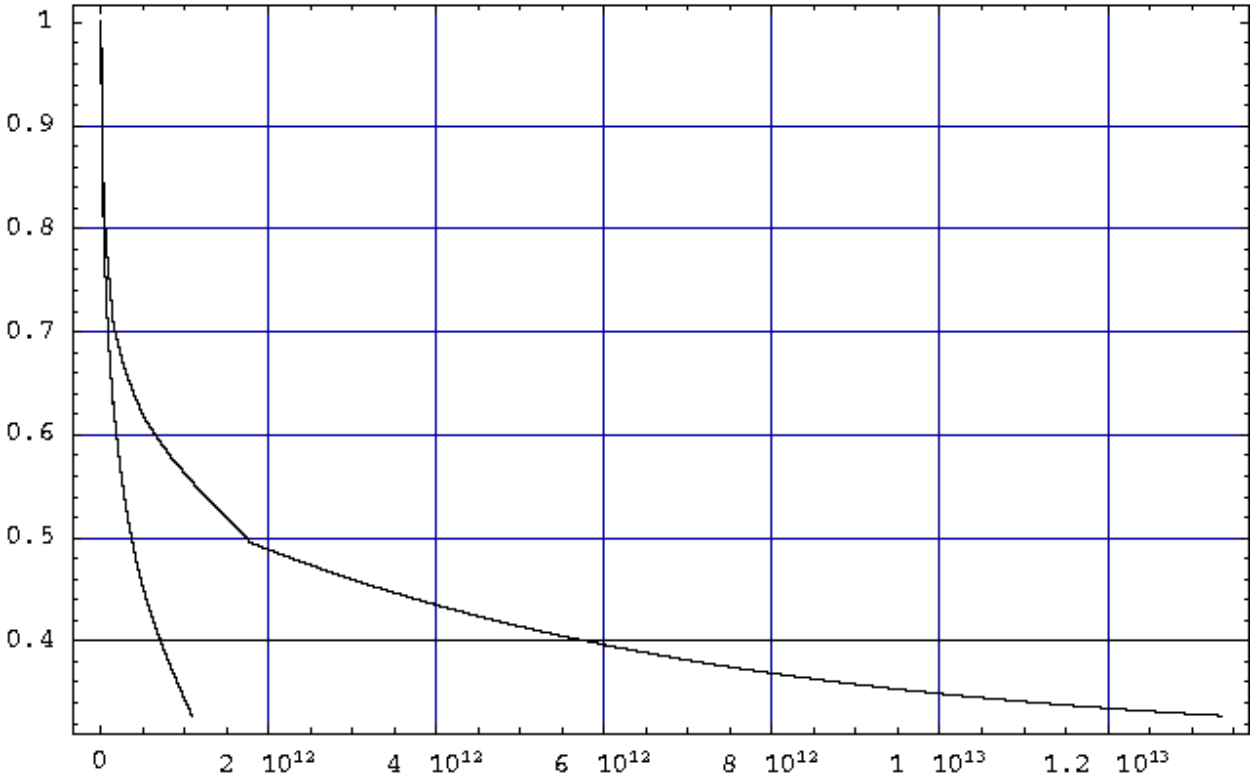

The energy expenditure figures on the X axis are in ergs/cm^3 of uncompressed uranium, the y axis gives the relative volume change (V/V_0). Shock compression, being less efficient, is the upper curve. It can be seen that as compression factors rise above 1.5 (a V/V_0 ratio of 0.67), the amount of work required for shock compression compared to isentropic compression rises

rapidly. The kink in the shock compression curve at V/V_0 of 0.5 is not a real phenomenon, it is due to the transition from experimental data to a theoretical Thomas-Fermi EOS.

It is interesting to note that to double the density of one cubic centimeter of uranium (18.9 grams) 1.7 x 10^12 ergs is required for shock compression. This is the amount of energy found in 40 grams of TNT, about twice the weight of the uranium. The efficiency of an implosion system at transferring high explosive energy to the core is generally not better than 30%, and may be worse (possibly much worse if the design is inefficient). This allows us the make a good estimate of the amount of explosive required to compress a given amount of uranium or plutonium to high density (a minimum of 6 times the mass of the fissile material for a compression factor of 2).

These curves also show that very high shock compressions (four and above) are so energetically expensive as to be infeasible. To achieve a factor of only 3, 7.1x10^11 ergs/g of uranium is required. Factoring implosion efficiency (30%), the high explosive (if it is TNT) must have a mass 56 times that of the material being compressed. Reports in the unclassified literature of compressions of four and higher can thus be safely discounted.

Compression figures for plutonium are classified above 30 kilobars, but there is every reason to believe that they are not much different from that of uranium. Although there are large density variations from element to element at low pressure, the low density elements are also the most compressible, so that at high pressures (several megabars) the plot of density vs atomic number becomes a fairly smooth function. This implies that what differences there may be in behavior between U and Pu at low pressure will tend to disappear in the high pressure region.

Actually, even in the low pressure region the available information shows that the difference in behavior isn't all that great, despite the astonishingly large number of phases (six) and bizarre behavior exhibited by plutonium at atmospheric pressure. The highest density phases of both metals have nearly identical atomic volumes at room pressure, and the number of phases of both metals drops rapidly with increasing pressure, with only two phases existing for both metals above 30 kilobars. The lowest density phase of plutonium, the delta phase, in particular disappears very rapidly. The amount of energy expended in compression at these low pressures is trivial. The compression data for uranium is thus a good substitute for plutonium, especially at high pressures and high compressions.

The shock and isentropic pressures required corresponding to the compression energy curves are shown below. The pressures shown on the X axis are in kilobars, the y axis gives the relative volume change (V/V_0).

Figure 4.1.6.2.1-2. Required Pressure for Shock and Isentropic Compression of Uranium

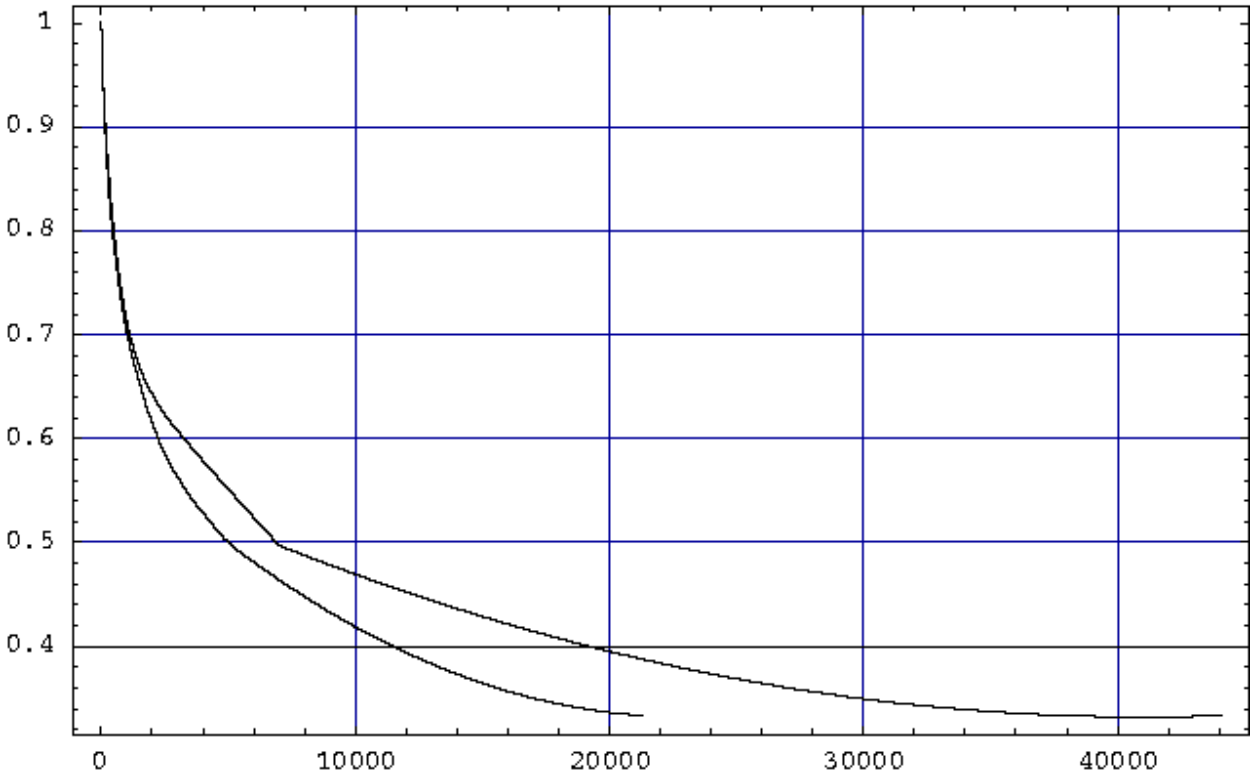

Since the compression energies of interest vary by many orders of magnitude over compressions ranging up to 3, it is often more convenient to look at logarithmic plots or energy. Figure 4.1.6.2.1-3, below, gives the isentropic curve from 10^7 ergs/cm^3 to 10^12 ergs/cm^3. Since the energy for shock compression is virtually identical to the isentropic value at small compressions, the curve for shock compression is given for compression energies of 10^10 erg/cm^3 (V/V_0 ~ 0.9)

Figure 4.1.6.2.1-3. Logarithmic Plot of Energy Required for Isentropic Compression of Uranium

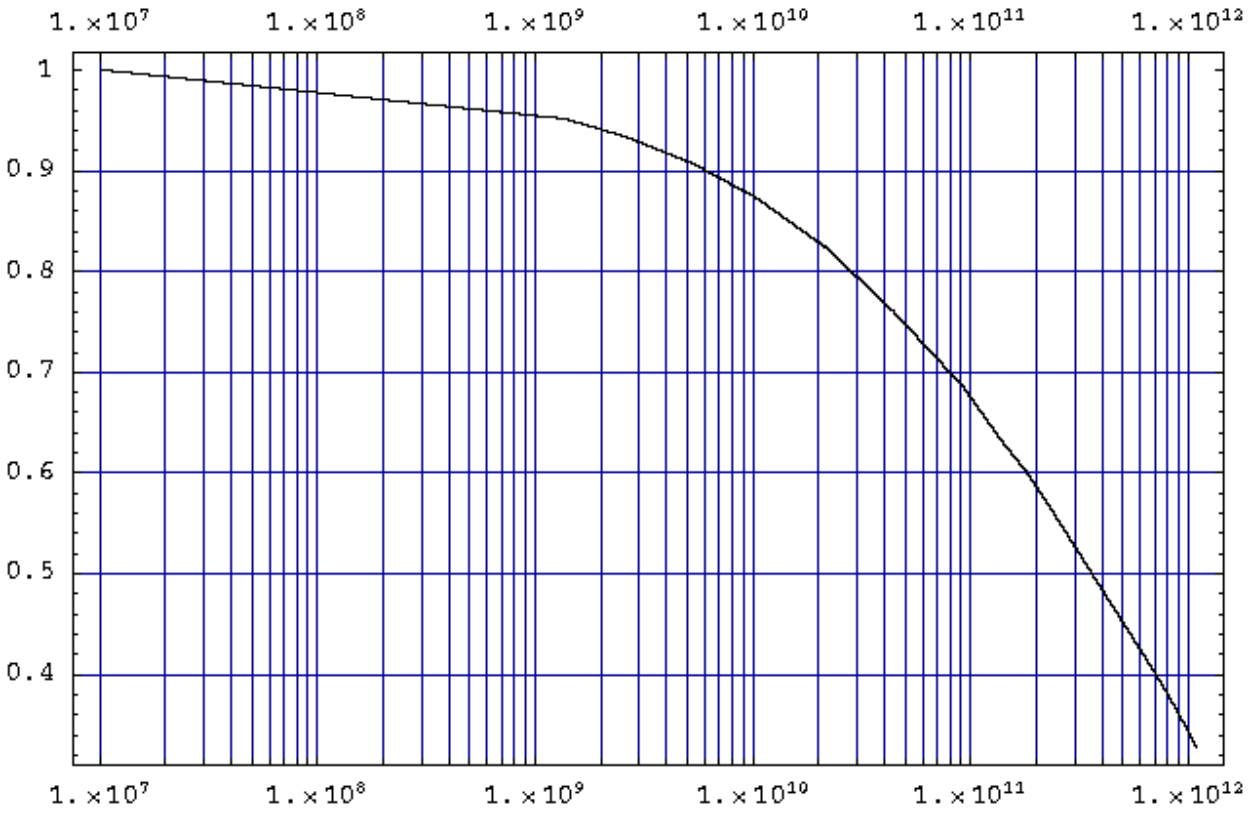

Figure 4.1.6.2.1-4. Logarithmic Plot of Energy Required for Shock Compression of Uranium

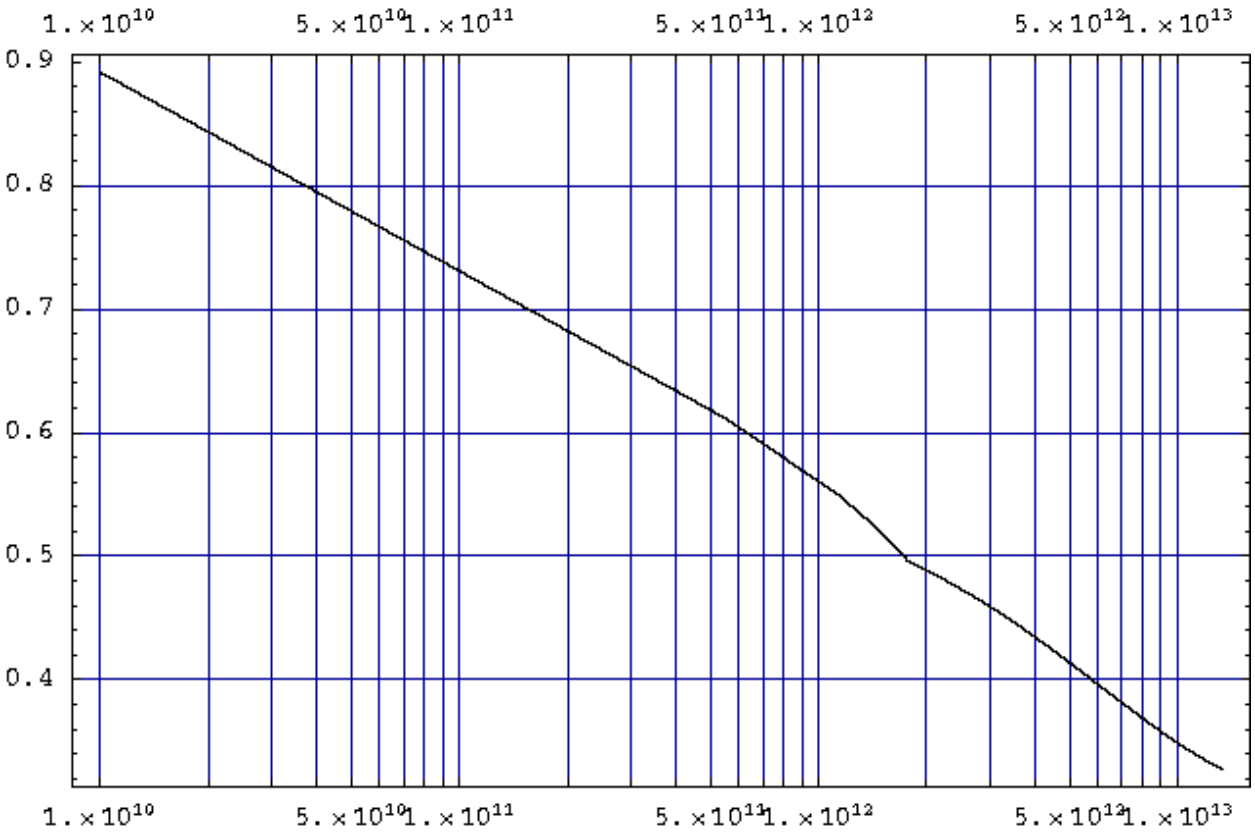

4.1.6.2.2 Shock Wave Generation Systems

The only practical means of generating shock waves in weapons is through the use of high explosives. When suitably initiated, these energetic materials support detonation waves: a self-sustaining shock wave that triggers energy releasing chemical reactions, and is driven by the expanding gases that are produced by these reactions.

Normally a high explosive is initiated at a single point. The detonation propagates as a convex detonation wave, with a more or less spherical surface, from that point.

To drive an implosion, a divergent detonation wave must be converted into a convergent one (or a planar one for linear implosion). Three approaches can be identified for doing this.

4.1.6.2.2.1 Multiple Initiation Points

In this approach, the high explosive is initiated simultaneously by a large number of detonators all over its surface. The idea is that if enough detonation points exist, then it will approximate the simultaneous initiation of the entire surface, producing an appropriately shaped shock from the outset.

The problem with this approach is that colliding shock waves do not tend to "smooth out", rather the reverse happens. A high pressure region forms at the intersection of the waves, leading to high velocity jets that outrun the detonation waves and disrupting the hoped for symmetry.

The multiple detonation point approach was the first one tried at Los Alamos during the Manhattan Project to build a spherical implosion bomb. Attempts were made to suppress the jetting phenomenon by constantly increasing the number of points, or by inserting inert spacers at the collision points to suppress the jets. The problems were not successfully worked out at the time.

Since the war this approach has been used with reasonable success in laboratory megagauss field experiments employing the simpler cylindrical geometry. There is also evidence of continuing US interest in this approach. It is not clear whether this technique has been successfully adapted for use in weapons.

4.1.6.2.2.2 Explosive Lenses

The basic idea here is to use the principle of refraction to shape a detonation wave, just as it is used in optics to shape a light wave.

Optical lenses use combinations of materials in which light travels at different speeds. This difference in speed gives rise to the refractive index, which bends the wave when it crosses the boundary between materials.

Explosive lenses use materials that transmit detonation or shock waves at different speeds. The original scheme used a hollow cone of an explosive with a high detonation velocity, and an inner cone of an explosive with a low velocity. The detonator initiates the high velocity explosive at the apex of the cone. A high velocity detonation wave then travels down the surface of the hollow cone, initiating the inner explosive as it goes by. The low velocity detonation wave lags behind, causing the formation of a concave (or planar) detonation wave.

With any given combination of explosives, the curvature of the wave produced is determined by the apex angle of the lens. The narrower the angle, the greater the curvature. However, for a given lens base area the narrower angle, the taller the lens, and the greater its volume. Both of these are undesirable in weapons, since volume and mass are at a premium.

To create a spherical implosion wave, a number of inward facing lenses need to be arranged on the surface of a sphere so that the convergent spherical segments that each produces merge into one wave. There is substantial advantage in using a large number of lenses. Having many lenses means that each lens has a small base area, and needs to produce a wave with a smaller curvature, both of which reduce the thickness of the lens layer. A more symmetrical implosion can probably be achieved with more lenses also.

It is important to have the lens detonation points (and optical axes) spaced as regularly as possible to minimize irregularities, and to make the height of each lens identical. The largest number of points that can be spaced equidistantly from their neighbors on the surface of a sphere is 20 - corresponding to the 20 triangular facets of an icosahedron (imagine the sphere encased in a circumscribed polyhedron, with each facet touching the sphere at one point). The next largest number is 12 - corresponding to the 12 pentagonal facets of the dodecahedron.

12 lenses, even 20 lenses, is an undesirably small number (although some implosion systems have used the 20 point icosahedral layout). A close approximation to strict regularity can be achieved with more points by interleaving a dodecahedron and icosahedron to produce a polyhedron tiled with hexagonal and pentagonal facets, 20 hexagons and 12 pentagons, for a total of 32 points. This pattern is the same familiar one found on a soccer ball, and was used as the original implosion system lens layout in Gadget, and other early US nuclear weapons.

Designs with 40, 60, 72, and 92 lenses have also been used (although these do not rely on Platonic solids for providing the layout pattern). More lenses lead to a thinner, less massive explosive lens shell, and greater implosion uniformity. The penalty for more lenses is more fabrication effort, and a more powerful and complex initiation system (not a trivial problem originally, but greatly simplified by modern pulse power technology). A simple implosion system could be very massive. The 32 point systems used in early US nuclear weapons had an external diameter of 1.4 m and weighed over 2000 kg. Current systems may be less than 30 cm, and weigh as little as 20 kg, but probably do not follow the same design approach as earlier weapons.

To a degree these multi-lens systems all suffer from the same shortcoming as the basic multi-point detonation approach: strict uniformity of the spherical implosion wave is unachievable. The detonation wave spreads out radially from each detonation point, so each wave produces a circular segment of a spherical wave. If you consider an icosahedron or a "soccer ball", you can see that when circles are inscribed in each of the regular polygons they touch each of their neighbor circles at one point. This marks the moment when the individual wavelets start to merge into a single wave. The gaps left between the inscribed circles however are irregular areas where distortions are bound to arise as the wave edges spread into them, possibly even leading to jetting.

Since the shock wave created by the lens exits from it at the velocity of the slow (and relatively weak) explosive, it desirable to have a layer of powerful explosive inside the lens system (perhaps the same one used as the fast lens component). This layer provides most of the driving force for the implosion, for the most part the lens system (which may well be much more massive) simply provides a mechanism for spherical initiation.

Ideally, the best combination of explosives is the fastest and slowest that are available. This provides the greatest possible refractive index, and thus bending effect, and allows using a wider lens angle. The fastest and slowest explosives generally known are HMX (octogen) and baratol respectively. HMX has a detonation velocity of 9110 m/sec (at a pressed density of 1.89), the dense explosive baratol (76% barium nitrate/24% TNT) has a velocity of 4870 m/sec (cast density 2.55). Explosives with slightly slower detonation velocities include the even denser plumbatol - 4850 m/sec (cast density 2.89) for a composition of 70% lead nitrate/30% TNT; and the relatively light boracitol - 4860 m/sec (cast density 1.55) for a composition of 60% boric acid/40% TNT. Mixtures of TNT with glass or plastic microspheres have proven to be an effective, light weight, and economical slow explosive in recent unclassified explosive lens work (I don't have data on their velocities though).

During WWII Los Alamos developed lenses using combination of Composition B (or Comp B) for the fast explosive (detonation velocity of 7920 m/sec, at a cast density 1.72), and baratol for the slow explosive.

Later systems have used the very fast HMX as a fast explosive, often as a plastic bonded mixture consisting almost entirely of HMX. Plumbatol, a denser and slightly slower explosive, may have been used in some later lens system designs. Boracitol is definitely known to have been used, probably in thermonuclear weapon triggers and perhaps in other types of weapons as well.

The idea of explosives lenses appears to have originated with M. J. Poole of the Explosives Research Committee in England. In 1942 he prepared a report describing a two-dimensional arrangement of explosives (RDX and baratol) to create a plane detonation wave. This idea was brought to Los Alamos in May 1944 by James Tuck, where he expanded it by suggesting a 3-D lens for creating a spherical implosion wave as a solution to making an implosion bomb. A practical lens design was proposed separately by Elizabeth Boggs of the US Explosives Research Laboratory, and by Johann Von Neumann. The Boggs proposal was the earlier of the two, although it was Von Neumann's proposal who directly influenced the Manhattan Project.

The task of developing a successful spherical implosion wave system is extremely difficult. Although the concept involved is simple, actually designing a lens is not trivial. The detonation wave velocity is affected by events occurring some distance behind the front. When the wave crosses from the fast explosive into the slow explosive it does not instantly assume the steady state detonation velocity of the slow explosive. Unlike the analogy with light, the velocity change is gradual and occurs over a significant distance. Since energy can be lost through the surface of the lens, thus reducing the fast wave velocity, the test environment of the lens also affects its performance. The behavior of a lens can only be calculated using sophisticated 2 and 3-D hydrodynamic computer codes that have been validated against experimental data.

Practical lens development generally requires a combination of experimentation, requiring precision explosive manufacture and sophisticated instruments to measure shock wave shape and arrival times, and numerical modelling (computer simulation) to extrapolate from test results. An iterative design, test, and redesign cycle allows the development of efficient, high-performance lenses.

During the Manhattan Project, due to the primitive state of computers and high explosive science and instrumentation, lenses could only be designed by trial and error (guided to some extent by scaling laws deduced from previous experiments). This required the detonation of over 20,000 test lens (and for each one tested, several were fabricated and rejected). When successful sub-scale implosion systems were scaled up to full size, it was discovered that the lenses had to be redesigned.

Assembling the lenses into a complete implosion system aggravates the design and development problems. To avoid shock wave collisions that disrupt symmetry, the surfaces of the lenses need to be aligned very accurately. In a spherical system, the implosion wave that is

created is completely hidden by the layer of detonating explosive. The chief region of interest is a small region in the center with perhaps < 0.1% the volume of the whole system. Very expensive diagnostic equipment and difficult experiments are required to study the implosion process, or even to verify that it works at all. Hemispherical tests can be quite useful though to validate lens systems before full spherical testing.

4.1.6.2.2.3 Advanced Wave Shaping Techniques

The conical lens design used by the Manhattan Project and early U.S. nuclear weapons is not the only lens design possible, or even the best. It had the crucial advantage of being simple in form (eliminating the need to design or fabricate complex shapes), and of having a single design variable - the cone apex angle. This made it possible to devise workable lenses with the crude methods then available. Other geometric arrangements of materials that transmit shocks slowly can be used to shape a convex shock into a concave one.

The shock slowing component of a lens, such as the inner cone of a conical explosive lens, does not really need to be another explosive. An inert substance that transmits a shock more slowly than the fast explosive detonation wave will also work. The great range of materials available that are *not* explosives gives much greater design flexibility. An additional (potential) advantage is that shock waves attenuate as they travel through non-explosive materials, and slow down. This can make lens design more complex, since this attenuation must be taken into account, but the reduced velocity can also lead to a more compact lens. Care must be taken though to insure that the attenuated shock remains strong enough to initiate the inner explosive layer.

By consulting the equation for shock velocity we can see that a high compressibility (low value of gamma) and a high density both lead to low shock wave velocities. An ideal material would be a highly compressible material of relatively high density. This describes an unusual class of filled plastic foams that have been developed at the Allied-Signal Kansas City Plant (the primary supplier of non-nuclear components for US nuclear weapons). It is quite possible that these foams were developed for use as wave shaping materials.

By extending the idea of custom tailoring the density and compressibility of materials, we can imagine that different arrangements of materials of varying properties can be used to reshape shock waves in a variety of ways.

Inserting low density materials, like solid or foam plastics, into explosives can also inhibit detonation propagation and allow the designer to "fold" the path the detonation wave must take. If suitable detonation inhibiting bodies are arranged in a grid inside a cone of high explosive, the same effect as the high explosive lens can be obtained with a lower lens density and with a larger apex angle.

French researchers have described advanced lens systems using alternating layers of explosive and inert material. This creates an anisotropic detonation velocity in the system, very slow across the layers, but fast along the them. A compact lens for producing spherically curved

waves has been demonstrated using a cylindrical version of this system, with a slow explosive between the inert layers, and a curved "nose cone"-like surface covered by fast explosive.

It is possible to completely and uniformly cover a sphere with circles if the number of lenses (and circles) is less than or equal to two. A single lens capable of bending a single detonation wave into a complete spherically convergent wave can, in principle, be made so that the resulting wave is entirely uniform. This extends the principle of the explosive lens to its most extreme form. It is also possible to use two lenses, each covering a hemisphere, which meet at the equator of the sphere and can smoothly join two hemispherical implosion waves.

The single point detonation system is illustrated below. This idea makes use of a cardioid-like logarithmic spiral:

```
      fffffff

    fsssssssssf

    fssssssssssssf

     fCsssssssssssssfD <- Detonator

    fssssssssssssf

    fssssssssssf          f = fast explosive

     fffffff              s = slow explosive

                          C = core
```

This not a very practical design as given. The thickness of the slow explosive on the detonator side would have to be considerable to achieve the necessary bending. Inserting detonation path folding spacers in the explosive could also dramatically reduce the size (but making manufacturing extremely difficult). A variation on this using the French layered explosive approach has also been proposed.

It is unlikely that a slow explosive would really be used for the inner slow lens component, since the velocity differential is not that great. The high degree of shock bending required strongly encourages using something that transmits shocks as slowly as possible such as an advanced inert material.

Such an implosion system would be extremely difficult to design and possibly to manufacture. The continuously varying 3-D surfaces would require considerable experimentation to perfect, and the surfaces would be a nightmare to machine. Once an acceptable shape were developed, and suitable molds or dies were made, the actual manufacture might be quite easy, requiring

only pressing of explosives and plastics into molds, or forming metal sheet in a die. The system would remain quite intolerant in any imperfections in dimensions or material however.

The difficulty in making compact and light implosion systems can be judged by the US progress in developing them. The initial Fat Man implosion system had a diameter of almost 60 inches. A significantly smaller system (30 inches) was not tested until 1951, a 22 inch system in mid-1952, and a 16 inch system in 1955. By 1955 a decade had passed since the invention of nuclear weapons, and hundreds of billions of dollars (in today's money) had been spent on developing and producing bombs and bomb delivery systems. These later systems must have used some advanced wave shaping technologies, which have remained highly classified. Clearly developing them is not an easy task (although the difficulty may be conceptual as much as technological).

4.1.6.2.2.4 Cylindrical and Planar Shock Techniques

Cylindrical and planar shock waves can be generated using the techniques previously described, making allowances for the geometry differences. A cylindrical shock can be created using the 2-D analog of the explosive lens, a wedge shaped lens with the same cross section as the conical version. A planar shock is simply a shaped shock with zero curvature.

A complete cylindrical implosion would require several parallel wedge-shaped explosive lenses arranged around the cylinder axis to form a star shape. To make the implosion truly cylindrical (as opposed to conical) it is necessary to detonate each of these lenses along the entire apex of the wedge simultaneously. This can be done by using a lens made out of sheets of high explosive (supported by a suitable backing) to create a plane shock. The edge of this sheet lens would join the apex of the wedge. This sheet lens need not extend out radially, it can join at an angle so that it folds into the space between the star points.

Some special techniques are also available based on the peculiar characteristics of the 1-D and 2-D geometries. The basic principle for these techniques is the "flying plate line charge", illustrated below.

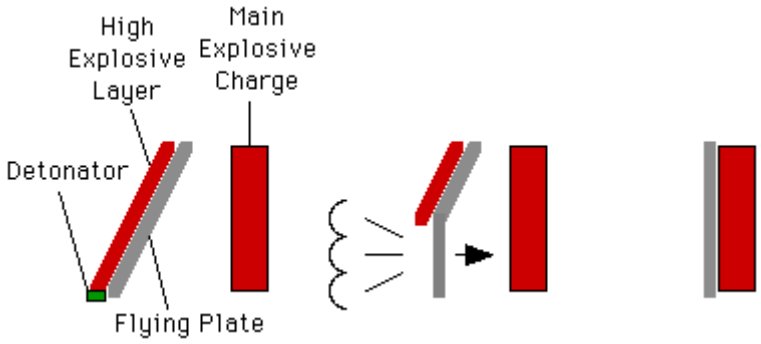

A metal plate is covered on one side with a sheet of explosive. It is detonated on one edge, and the detonation wave travels across the plate. As it does so the detonation accelerates the plate, driving it to the right. After the explosive has completely detonated the flying plate will be flat again. The angle between the original stationary plate and the flying plate is determined by the ratio between the detonation velocity, and the velocity of the accelerated plate. When this high

velocity plate strikes the secondary explosive charge the shock will detonate it, creating a planar detonation.

As described above, the system doesn't quite work. A single detonator will actually create a circular detonation front in the explosive sheet, expanding from the initiation point. This can be overcome by first using a long, narrow flying plate (a flying strip if you will) to detonate the edge of wide plate. This wide plate can then be used to initiate the planar detonation.

The flying strip approach can also be used to detonate the cylindrical lens system described above in place of the sheet lens.

The flying plate scheme can be easily extended to create cylindrical detonations.

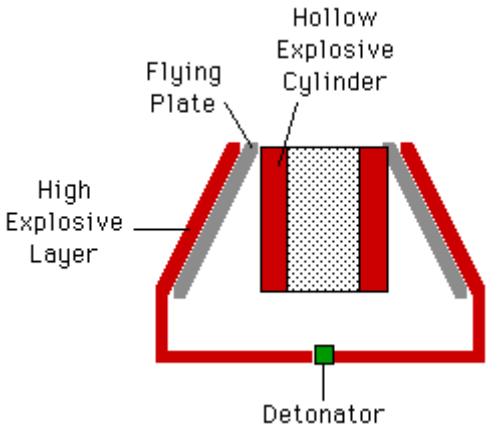

This is a cross section view of a hollow truncated cone covered by a layer of explosives. The wide end of the cone is joined to a sheet of explosives with a detonator in the center.

The single detonator located on the axis causes an expanding circular detonation in the explosive sheet. When the shock wave reaches the perimeter, it continues travelling along the surface of the cone. The cone collapses starting at the wide end. The angle of the cone is such that a cylindrical flying plate is created that initiates a cylindrical detonation in the secondary explosive.

Flying plate systems are much easier to develop than explosive lenses. Instrumentation for observing their behavior is relatively simple. Multiple contact pins and an oscilloscope can easily measure plate motion, and well established spark gap photography can image the plate effectively.

4.1.6.2.2.5 Explosives

The choice of explosives in an implosion system is driven by the desire for high performance, safety, ease of fabrication, or sometimes by special properties like the slow detonation velocity needed in explosive lenses.

The desire for high performance leads to the selection of very energetic explosives that have very high detonation velocities and pressures (these three things are closely correlated). The highest performance commonly known explosive is HMX. Using HMX as the main explosive will provide the greatest compression. HMX was widely used in US weapons from the late fifties on into the 1970s, often in a formula called PBX-9404 (although this particular formulation proved to have particularly serious safety problems - causing eight fatalities in a six month period in 1959 among personnel fabricating the explosive). HMX is known to be the principal explosive in many Soviet weapon designs since Russia is selling the explosive extracted from decommissioned warheads for commercial use. The chemically related RDX is a close second in power. It was the principal explosive used in most early US designs, in the form of a castable mixture called Composition B.

In recent years the US has become increasingly concerned with weapon safety, following some prominent accidents in which HE detonation caused widespread plutonium contamination and in the wake of repeated fatal explosions during fabrication. Many of the high energy explosives used, such as RDX and HMX, are rather sensitive to shock and heat. While normally an impact on the order of 100 ft/sec is required to detonate one these explosives, if a sliding or friction-producing impact occurs then these explosives can be set off by an impact as slow as 10 ft/sec (this requires only a drop of 18 inches)! This has led to the use of explosives that are insensitive to shock or fire. Insensitive explosives are all based on TATB, the chemical cousin DATB lacks this marked insensitivity. These explosives have very unusual reaction rate properties that make them extremely insensitive to shock, impact, or heat. TATB is reasonably powerful, being only a little less powerful than Comp B. A composition known as PBX-9504 has been developed that adds 15% HMX to a TATB mixture, creating a compromise between added power and added sensitivity.

Another very strong explosive called PETN has not been used much (or at all) as a main explosive in nuclear weapons due to its sensitivity, although it used in detonators.

Fabricating explosives for implosion systems is a demanding task, requiring rigid quality control. Many explosive components have complex shapes, most require tight dimensional tolerances, and all require a highly uniform product. Velocity variations cannot be greater than a few percent. Achieving such uniformity means carefully controlling such factors as composition, purity, particle size, crystal structure, curing time and curing temperature.

Casting was the first method used for manufacturing implosion components since a very homogenous product can be produced in fairly complex shapes. Unfortunately the most desirable explosives do not melt, which makes casting of the pure explosive impossible. The original solution adopted by the US to this problem was to use castable mixtures of the desired explosive and TNT. TNT is the natural choice for this, being the only reasonably powerful, easily melted explosive available. Composition B, the first explosive used, typically consisted of 63% RDX, 36% TNT, and 1% wax (cyclotol, a mixture with a higher proportion of RDX to TNT, was used later). Great care must be taken to ensure that the slurry of solid explosive and melted TNT is uniform since settling occurs. Considerable attention must be paid to controlling the particle size of the solid explosive, and to monitoring the casting, cooling, and curing processes.

Mold making is also a challenging task, requiring considerable experimentation at Los Alamos before an acceptable product could be made.

Pressing is a traditional way of manufacturing explosives products, but its inability to make complex shapes, and problems with density variations and voids prevented its use during WWII. Plastic explosives (that is - soft, pliable explosives) can be pressed into uniform complex shapes quite easily, but their lack of strength make them unattractive in practical weapon designs.

During the forties and fifties advances in polymer technology led to the creation of PBXs (plastic bonded explosives). These explosives use a polymer binder that sets during or after fabrication to make a rigid mass. The first PBX was developed at Los Alamos in 1947, an RDX-polystyrene formulation later designated PBX 9205. Some early work used epoxy binders that harden after fabrication through chemical reactions, but current plastic binders are thermosetting resins (possibly in combination with a plasticizer). Explosive granules are coated with the plastic binder and formed by pressing, usually followed by machining of the billet.

The desire for maximum explosive energy has led to the selection of polymers and plasticizers that actively participate in the explosion, releasing energy through chemical reactions. Emphasis on this has led to undesirable side effects - like sensitization of the main explosive (as occurred with PBX-9404), or poor stability. In the 1970s the W-68 warhead, the comprising large part of the U.S. submarine warhead inventory, developed problems due to decomposition of the LX-09 PBX being used, requiring the rebuilding of 3,200 warheads. LX-09 also exhibited sensitivity problems similar to PBX-9404, in 1977 three men were killed at the Pantex plant in Amarillo from a LX-09 billet explosion.

Normally the explosive and polymer binder are processed together to form a granulated material called a molding powder. This powder is formed using hot pressing - either isostatic (hydrostatic) or hydraulic presses, using evaluated molds (1 mm pressure is typical). The formed material may represent the final component, but normally additional machining to final specifications is required.

PBXs contain a higher proportion of the desired explosive, possess greater structural strength, and also don't melt. These last two properties make them easier to machine to final dimensions. Plastic bonding is very important in insensitive high explosives (IHEs), since mixing the insensitive explosives with the more sensitive TNT would defeat the purpose of using them.

PBX was first used in a full-scale nuclear detonation during the Redwing Blackfoot shot in June 1956. PBXs have replaced melt castable explosives in all US weapons. The PBX compositions that have been used by the U.S. include PBX-9404, PBX-9010, PBX-9011, PBX-9501, LX-04, LX-07, LX-09, LX-10, LX-11. Insensitive PBXs used are PBX-9502 and LX-17.

Explosive Compositions Used In U.S. Nuclear Weapons

- Baratol: 76% barium nitrate, 24% TNT (typical)
  Low velocity castable explosive used in early explosive lenses.

- Boracitol: 60% boric acid, 40% TNT (typical)
  Low velocity castable explosive used in later explosive lens designs.
- Composition B: 63% RDX, 36% TNT, 1% wax (typical)
  High velocity castable main explosive used in early nuclear weapons (e.g. Fat Man; Mks 4, 5 and 6), also MK 28 and MK 53 (latter warhead still in service).
- Cyclotol: 75% RDX, 25% TNT
  High velocity castable main explosive, basically just Comp B with a higher RDX content for higher performance. Used in MK 28 and MK-53 (latter warhead still in service). Substituted for PBX-9404 when unacceptable sensitivity problems arose.
- LX-04: 85% HMX, 15% Viton A
  High velocity PBX main explosive. Used in W-62 and W-70.
- LX-07: 90% HMX, 10% Viton A
  High velocity PBX main explosive. Used in W-71.
- LX-09: 93% HMX, 4.6% pDNPA, 2.4% FEFO
  High velocity PBX. Main explosive used in the W-68 warhead. Withdrawn from use due to aging problems (binder/plasticizer exudation). Serious safety problems.
- LX-10: 95% HMX, 5% Viton A; and LX-10-1: 94.5% HMX, 5.5% Viton A
  High velocity PBX main explosive. Replaced LX-09 in W-68. Also used in W-70; W-79; and W-82.
- LX-11: 80% HMX, 20% Viton A
  High velocity PBX main explosive. Used in W-71.
- LX-17: 92.5% TATB, 7.5% Kel-F 800
  High velocity insensitive PBX. One of two IHEs in use. Used in B-83; W-84; W-87; and W-89. Stockpile-monitoring of the W87 warhead shows some evidence of stiffening with age, perhaps due to an increase in the crystallinity of the binder.
- PBX-9010: 90% RDX, 10% Kel-F
  High velocity PBX main explosive. Used in MK 43 and W-50.
- PBX-9011: 90% HMX, 10% Estane
  High velocity PBX main explosive. Used in MK 57 Mods 1 and 2.
- PBX-9404: 94% HMX, 3% NC, 3% CEF
  High velocity PBX main explosive. Widely used - MK 43; W-48; W-50; W-55; W-56; MK 57 Mod 2; MK/B 61 Mods 0, 1, 2, 5; and W-69. Serious safety problems.
- PBX-9501: 95% HMX, 2.5% Estane, 2.5% BDNPA-F
  High velocity PBX main explosive. Used in W-76; W-78; and W-88.
- PBX-9502: 95% TATB, 5% Kel-F
  High velocity insensitive PBX. Principal IHE in recent US weapon designs, currently being backfitted to earlier warheads replace other plastic bonded explosives. Used in B-61 Mods 3, 4, 6-10; W-61; W-80; W-85; W-90; and W-91.
- Plumbatol: 70% lead nitrate, 30% TNT (typical)
  The use of this low velocity castable explosive in US nuclear weapons is speculative.

Explosives And Binder Ingredients Used In U.S. Nuclear Weapons

- Barium nitrate: Heavy metal oxidizer used in baratol slow explosive mixture.

- BDNPA-F: Liquid polymer/plasticizer mixture used in PBX compositions. 50% bis(2,2-dinitropropyl), 50% acetal/bis(2,2-dinitropropyl)formal (plasticizer)
- Boric Acid: Low density, low atomic number inert material used in boracitol slow explosive mixture.
- CEF: Plasticizer used in PBX mixtures. tris-beta-chloroethylphosphate.
- DATB: Main high explosive, insensitive. 2,4,6-trinitro-1,3-benzenediamine; also called DATNB, diamino trinitrobenzene.
- DNPA (pDNPA): Solid explosive used in a binder mixture. 2,2-dinitropropyl acrylate.
- FEFO: Liquid explosive used in a binder mixture. 1,1-[methylenebis(oxy)]-bis-[2- fluoro-2,2-dinitroethane].
- HMX: Main high explosive, very powerful. Octahydro-1,3,5,7-tetranitro-1,3,5,7-tetrazocine; also called beta-HMX, octogen, cyclotetramethylene tetranitramine, (HMX is WWII code name, from His Majesty's eXplosive). Dual use material, export restricted.
- HNS: Relatively insensitive, a very heat stable high explosive, used in slapper detonators. 1,1'-(1,2-ethylenediyl) bis-(2,4,6-trinitrobenzene); also called hexanitrostilbene. Dual use material, export restricted.
- Kel-F: Inert plastic binder. Copolymer consisting of chlorotrifluoroethylene / vinylidine fluoride (3:1 ratio).
- Lead nitrate Heavy metal oxidizer used in plumbatol slow explosive mixture.
- NC: Solid explosive used as a plastic binder. Nitrocellulose.
- PETN: Sensitive powerful high explosive, used in detonators. 2,2- bis[(nitroxy)methyl]-1,3-propanediol dinitrate; also called pentaerythritol tetranitrate.
- RDX: Main high explosive, powerful. hexahydro-1,3,5-trinitro-1,3,5-triazine; also called cyclonite, hexogen. Dual use material, export restricted.
- TATB: Main high explosive, very insensitive and heat stable. Special fine-grained TATB used in boosters. 2,4,6-trinitro-1,3,5-benzenediamine; also called TATNB, triaminotrinitrobenzene. Dual use material, export restricted. Produced on an industrial scale in the U.S. at a cost of $90 to $250/kg. Currently available to customers outside DOE for about $200/kg.
- TNT: Main high explosive, used as a meltable binder. 2-methyl-1,3,5- trinitrobenzene; also called trinitrotoluene.
- Viton A: Rubbery solid used as a plastic binder. Copolymer consisting of 60% Vinylidine fluoride/40% hexafluoropropylene.

Table 4.1.6.2.2.5-1. Basic Properties Of Explosives Used In Us Nuclear Weapons

EXPLOSIVE	DETONATION VELOCITY m/sec	DETONATION PRESSURE kilobars	DENSITY	SENSITIVITY
HMX	9110	390	1.89/pressed	Moderate
LX-10	8820	375	1.86/pressed	Moderate

LX-09	8810	377	1.84/pressed	Moderate
PBX-9404	8800	375	1.84/pressed	Moderate
RDX	8700	338	1.77/pressed	Moderate
PETN	8260	335	1.76/pressed	High
Cyclotol	8035	-	1.71/cast	Low
Comp B 63/36	7920	295	1.72/cast	Low
TATB	7760	291	1.88/pressed	Very Low
PBX-9502	7720	-	1.90/pressed	Very Low
DATB	7520	259	1.79/pressed	Low
HNS	7000	200	1.70/pressed	Low
TNT	6640	210	1.56/cast	Low
Baratol 76/24	4870	140	2.55/cast	Moderate
Boracitol 60/40	4860	-	1.55/cast	Low
Plumbatol 70/30	4850	-	2.89/cast	Moderate

4.1.6.2.2.6 Detonation Systems

Creating a symmetric implosion wave requires close synchronization in firing the detonators. Tolerances on the order of 100 nanoseconds are required.

Conventional detonators rely on electrically heating a wire, which causes a small quantity of a sensitive primary explosive to detonate (lead azide, mercury fulminate, etc.). The primary usually then initiates a secondary explosive, like PETN or tetryl, which fires the main charge.

The process of resistively heating the wire, followed by heat conduction to the primary explosive until it reaches detonation temperature requires a few milliseconds, with correspondingly large timing errors. Conventional detonators thus lack the necessary precision for firing an implosion system.

One approach to reducing the duration of action of the detonator is to send a sudden, powerful surge of current through a very fine wire (made of gold or platinum), heating it to the point of vaporization. This technique, called an exploding wire or exploding bridge wire (EBW) detonator, was invented by Luis Alvarez at Los Alamos during the Manhattan Project. Current

surge rise times of a fraction of a microsecond are feasible, with a spread in detonation times of a few nanoseconds.

An exploding wire detonator can be used to initiate a primary explosive (usually lead azide), as in a conventional detonator. But if the current surge is energetic enough, then the exploding wire can directly initiate a less sensitive booster explosive (usually PETN). The advantage of doing this is that the detonation system is extremely safe from accidental activation by heat, stray currents, or static electricity. Only very powerful, very fast current surges can fire the detonators. This type of exploding wire detonator is one of the safest types of detonators known. The disadvantage is the need to supply those very powerful, very fast current surges. A typical EBW requires 5 KV, with a peak current of at least 500-1000 amps. A few kiloamps is more typical of most EBW detonators, but a multi-EBW system would probably try to minimize the required current. With sufficient care in detonator design and construction, inherent detonator accuracies of better than 10 nanoseconds are achievable.

Since WWII, a number of detonator designs based on exploding foils have been developed. Exploding foil detonators could be used to fire the booster explosive directly, as in EBW detonators, but generally this implies the use of different concept called a "slapper" detonator. This idea (developed at Lawrence Livermore) uses the expanding foil plasma to drive another thin foil or plastic film to high velocities, which initiates the explosive by impacting the surface. Normally the driving energy is provided entirely by heating of the foil plasma from the current passing through it, but more sophisticated designs may use a "back strap" to create a magnetic field that drives the plasma forward. Slappers are fairly efficient at converting electrical energy into flyer kinetic energy, it is not hard to achieve 25-30% energy transfer.

A typical slapper detonator consists of an explosive pellet pressed to a high density for maximum strength (plastic bonded explosives can also be used). Next to the explosive pellet is an insulation disk with a hole in the center which is set against the explosive pellet. An insulating "flyer" film, such as Kapton or Mylar with a metal foil etched to one side is placed against the disk. A necked down section of the etched foil acts as the bridgewire. The high current firing pulse causes vaporization of the necked down section of the foil. This then shears the insulated flyer which accelerates down the barrel of the disk and impacts the explosive pellet. This impact energy transmits a shock wave into the explosive causing it to detonate.

Another possible advantage of a slapper detonator is the ability to initiate an area of explosive surface rather than a point. This may make compact implosion systems easier to design.

This system has several advantages over the EBW detonator. These include:

1. the metal bridge is completely separated from the explosive by an insulating film and an air gap (the bridgewire of an EBW is in direct contact with the explosive);
2. the explosive can be packed to a high (near crystal) density;
3. the energy requirement to fire the detonator is lower; and
4. very insensitive explosives such as HNS can be detonated, which is extremely difficult with the EBW approach.

Exploding wire detonators were used in the first atomic device, but have since been replaced in the U.S. arsenal by foil slappers, and very probably in all other arsenals as well. Due to the ability of slapper detonators to use insensitive primary explosives, these are almost certainly used with all insensitive high explosive equipped warheads (unless supplanted by an even more advanced technology - like laser detonators).

More recently laser detonating systems have been developed. These use a high power solid state laser to deliver sufficient energy in the form of a short optical pulse to initiate a primary or booster explosive. The laser energy is conducted to the detonator by a fiber optic cable. This is a safe detonator system, but the laser and its power supply is relatively heavy. A typical system might use a 1 W solid state laser to fire a single detonator. It is not known if this system has been used in any nuclear weapons.

Another fast detonator is the spark gap detonator. This uses a high voltage (approx. 5 KV) spark across a narrow gap to initiate the primary explosive. If a suitably sensitive primary explosive is used (lead azide, or the especially sensitive lead styphnate) then the current required is quite small, and a modest capacitor can supply sufficient power (10-100 millijoules per detonator). The chief disadvantage of this detonator design is that it is one of the least safe known. Static charges, or other induced currents, can very easily fire a spark gap detonator. For this reason they have probably never been used in deployed nuclear weapons.

Detonation systems require a reasonably compact and light high speed pulse power supply. To achieve accurate timing and fast response requires a powerful power source capable of extremely fast discharge, as well as fast, accurate, and reliable switching components, and close attention to managing the inductance of the entire system.

The normal method of providing the power for an EBW multi-detonator system is to discharge a high capacitance, high voltage, low inductance capacitor. Voltage range is several kilovolts, 5 KV is typical. Silicone oil filled capacitors using Kraft paper, polypropylene, or Mylar dielectrics are suitable types, as are ceramic-type capacitors. Compact power supplies for charging capacitors are readily available.

The capacitor must be matched with a switch that can handle high voltages and currents, and transition from a safe non-conducting state to a fully conducting one rapidly without adding undue inductance to the circuit. A variety of technologies are available: triggered spark gaps, krytrons, thyratrons, and explosive switches are some that could be used.

The current rise time of the firing pulse can actually be much longer than the required timing accuracy since the firing of an EBW detonator is basically determined by achieving a threshold current. As long as the current rise is synchronous for all detonators, they will fire simultaneously. Still a rise time of no more than 2-3 microseconds is desirable.

The capacitance required for a 5 KV EBW is on the order of 1 microfarad per detonator. A 32 detonator system (like Fat Man) thus requires at least 32 mF and to produce a 32 kA current surge. For a rise time of 3 microseconds this requires no more than 100 nanohenries of total

inductance. A modern plastic cased capacitor of 40 microfarads, rated at 5 KV, with 100 nanohenries of inductance weighs about 4 kg.

Triggered spark gaps are sealed devices filled with high pressure air, argon, or SF6. A non-conducting gap between electrodes is closed by applying a triggering potential to a wire or grid in the gap. Compact versions of these devices are typically rated at 20-100 KV, and 50-150 kiloamps. The triggering potential is typically one-half to one-third the maximum voltage, with switch current rise times of 10-100 nanoseconds.

Krytrons are a type of cold cathode trigger discharge tube. Krytrons are small gas filled tubes. Some contain a small quantity of Ni-63, weak beta emitter (92 yr half-life, 63 KeV) that keeps the gas in a slightly ionized state. Applying a trigger voltage causes an ionization cascade to close the switch. These devices have maximum voltage ratings from 3 to 10 KV, but peak current rating of only 300-3000 amps making them unsuitable for directly firing multiple EBW detonators. They are small (2 cm long), rugged, and accurate (jitter 20-40 nanoseconds) however, and are triggered by voltages of only 200-300 V. They are very convenient then for triggering other high current devices, like spark gaps, by discharging through a pulse current transformer (they can, in turn, be conveniently triggered using a small capacitor, pulse transformer, and a thyristor). Krytrons are used commercially in powerful xenon flash lamp systems, among other uses. Krytrons have faster response times than other types of trigger discharge tubes. A vacuum tube relative of the krytron, the sprytron, is very similar and has very high radiation resistance. It is probably the sprytron that is actually used in U.S. nuclear weapons. The only manufacturer of krytrons and sprytrons is EG&G, the same company that provided the spark gap cascades for Gadget, Fat Man, and other early atomic weapons.

Other switching techniques that have been developed are explosive switches, and various other vacuum or gas-filled tube devices like hydrogen thyratrons and arc discharge tubes. An explosive switch uses the shock wave from an explosive charge to break down a dielectric layer between metal plates. Both this technique and the thyratron were under development at Los Alamos at the end of WWII.

Detonators are wired in parallel for reliability and to minimize inductance. For additional reliability, redundant detonation circuits may be used. In the Fat Man bomb the detonators were wired in parallel in spark gap triggered circuits. There were four detonating circuits, any two of which provided sufficient power for all 32 detonators. Each detonator was wired to two different circuits so that the failure of any one detonator circuit (and up to two of them) would not have affected the implosion. The whole system was fired by a spark gap cascade - the trigger spark gap supplied a current surge to fire the four main circuits simultaneously.

With sufficient care timing accuracies of 10 nanoseconds are achievable, which is probably better than practical implosion systems require (100 nanosecond accuracy is more typical).

Although the types of switches and capacitors mentioned here are, for the most part, available from many commercial sources and have many commercial uses, they are nonetheless subject

to dual use export controls. Attempts to export of krytrons illegally has been especially well publicized over the years, but they are not the only such devices suitable for these applications.

The detonator bridge wire used in EBWs is typically made of high purity gold or platinum, 20-50 microns wide and about 1 mm long. PETN is invariably used as the explosive, possibly with a tetryl booster charge. Slapper detonators use metal foils (usually aluminum, but gold foil would work well also) deposited on a thin plastic film (usually Kapton). A wider variety of primary explosives can be used. PETN or HMX may have been used in slappers used in earlier weapon systems, but weapons using IHE probably use the highly heat stable HNS.

A possible substitute for a capacitor bank in a detonation system is an explosive generator, also called a flux compression generator (FCG). This consists of a primary coil that is energized to create a strong magnetic field by a capacitor discharge. At the moment of maximum field strength an explosive charge drives a conducting plate into the field, rapidly compressing it. The rising magnetic field induces a powerful high voltage current in a secondary coil. Any of the switching technologies mentioned above can then be used to switch the load to the detonating system. A substantial fraction of the chemical energy of the explosive can be converted to electrical power in this way.

FCGs can potentially provide ample power for detonators and external neutron initiators at a very modest weight. Extensive research on these generators has been conducted at Los Alamos and Lawrence Livermore, and they are known to have been incorporated into actual weapon designs (possibly the Mk12, which had 92 initiation points).

4.1.6.2.3 Implosion Hardware Designs

Once created implosion shocks can be used to drive different implosion hardware systems. By implosion hardware, I mean systems of materials that are inert from the viewpoint of chemical energy release: the fissile material itself, and any reflectors, tampers, pushers, drivers, buffers, etc.

One approach to designing an implosion hardware system is to simply use the direct compression of the explosive generated shock wave to accomplish the desired reactivity insertion. This is the "solid pit design" used in Gadget and Fat Man.

A variety of other designs make use of high velocity collisions to generate the compressive shocks for reactivity insertion. These velocities of course are obtained from the energy provided by the high explosive shocks.

4.1.6.2.3.1 Solid Pit Designs

Since shock waves inherently compress the material through which they pass, an obvious way of using the implosion wave is simply to let it pass through the fissile core, compressing it as it converges on the center. This technique can (and has) been used successfully, but it has some inherent problems not all of which can be remedied.

First, the detonation pressure of available explosives (limit 400 kilobars) is not high enough for much compression. A 25% density increase is all that can be obtained in uranium at this pressure, delta-phase plutonium can reach 50% due to the low pressure delta->alpha phase transformation. This pressure can be augmented in two ways: by reflecting the shock at high impedance interfaces, and by convergence.

Since the fissile material is about an order of magnitude denser than the explosive itself, the first phenomenon is certain to occur to some extent. It can be augmented by inserting one or more layers of materials of increasing density between the explosive and the dense tamper and fissile material in the center. As a limit, shock pressure can double when reflected at an interface. To approach this limit the density increase must be large, which means that no more than 2 or 3 intermediate layers can be used.

The second phenomenon, shock convergence, is limited by the ratio of the fissile core radius to the outer radius of the implosion hardware. The intensification is approximately proportional to this ratio. A large intensification thus implies a large diameter system - which is bulky and heavy.

Another problem with the solid pit design is the existence of the Taylor wave, the sharp drop in pressure with increasing distance behind the detonation front. This creates a ramp-shaped shock profile: a sudden jump to the peak shock pressure, followed by a slope down to zero pressure a short distance behind the shock front. Shock convergence actually steepens the Taylor wave since the front is augmented by convergence to a greater degree than the material behind the front (which is at a larger radius). If the Taylor wave is not suppressed, by the time the shock reaches the center of the fissile mass, the outer portions may have already expanded back to their original density.

The use of intermediate density "pusher" layers between the explosive and the tamper helps suppress or flatten the Taylor wave. The reflected high pressure shock reinforces the pressure behind the shock front so that instead of declining to zero pressure, it declines to a pressure equal to the pressure jump at the reflection interface. That is, if P is the initial shock pressure, and P -> 0 indicates a drop from P to zero through the Taylor wave, then the reflection augments both by p:(P + p) -> (0 + p).

The Gadget/Fat Man design had an intermediate aluminum pusher between the explosive and the uranium tamper, and had a convergence factor of about 5. As a rough estimate, one can conclude that the 300 kilobar pressure of Composition B could be augmented by a factor of 4 by shock reflection (doubling at the HE/Al interface, and the Al/U interface), and a factor of 5 by convergence, leading to a shock pressure of 6 megabars at the plutonium core. Assuming an alpha phase plutonium equation of state similar to that of uranium this leads to a compression of a bit less than 2, which when combined with the phase transformation from delta to alpha gives a maximum density increase of about 2.5. The effective compression may have been significantly less than this, but it is generally consistent with the observed yield of the devices.

4.1.6.2.3.2 Levitated Core Designs

In the solid pit design, the Taylor wave is reduced but not eliminated. Also, the kinetic energy imparted by the convergent shock is not efficiently utilized. It would be preferable to achieve uniform compression throughout the fissile core and tamper, and to be able to make use of the full kinetic energy in compressing the material (bringing the inward motion of material in the core to a halt at the moment of maximum compression).

This can be accomplished by using a shell, or hollow core, instead of a solid one (see Section 3.7.4 Collapsing Shells). The shell usually consists of an outer layer of tamper material, and an inner layer of fissile material. When the implosion wave arrives at the inner surface of the shell, the pressure drops to zero and an unloading wave is created. The shock compressed material (which has also been accelerated inward) expands inward to zero pressure, converting the compression energy into even greater inward directed motion (approximately doubling it). In this way energy loss by the outward expansion of material in the Taylor wave region is minimized.

Simply allowing this fast imploding hollow shell to collapse completely would achieve substantial compression. In practice this is never done. It is more efficient to allow the collapsing shell to collide with a motionless body in the center (the "levitated core"), the collision creating two shock waves - one moving inward to the center of the stationary levitated core (accelerating it inward), and one moving outward through the imploding shell (decelerating it). The pressure between these two shocks is initially constant so that when the converging shock reaches the center of the core, the region extending from the center out to the location of expanding shock has achieved reasonably even and efficient compression.

I use word "reasonably" because the picture is a bit more complicated than just described. First, by the time the shell impacts the levitated core it has acquired the character of a thick collapsing shell. The inner surface will be moving faster than the outer surface, and a region close to the inner surface will be somewhat compressed. Second, the inward and outward moving shocks do not move at constant speed. The inward moving shock is a classical converging shock with a shock velocity that accelerates and strengthens all the way to the center. The outward moving shock is a diverging or expanding shock that slows down and weakens.

In the classical converging shock region (the levitated core, and the innermost layer of the colliding shell) high compression is achieved and the material is brought to a halt when the shock reaches the center. In the outer diverging region, only about half of the implosion velocity is lost when the diverging shock compresses and decelerates it, and there is insufficient time for inward flow to bring it to a halt before the converging shock reaches the center. Thus the outer region is still collapsing (slowly) when the inner shock reaches complete convergence (assuming that the outer shock has not yet reached the surface of the pit (tamper shell plus core) and initiated an inward moving release wave).

Immediately after the converging shock reaches the center, the shock rebound begins. This is an outward moving shock that accelerates material away from the center, creating an expanding low density region surrounded by a layer compressed to an even greater degree than in the initial implosion. Once the rebound shock expands to a given radius the average density of the volume within that radius falls rapidly.

For a radius well outside the classical converging shock region, the true average density may continue to increase due to the continuing collapse of the outer regions until the rebound shock arrives. The structure of the shell/core system at the time of rebound shock arrival is actually hollow - a low density region in the center with a highly compressed shell, but the average density is at a maximum. Whether this configuration is acceptable or not depends on the weapon design, it may be acceptable in a homogenous un-boosted core but will not be acceptable in a boosted or a composite core design where high density at the center is desired.

Since the divergence of the outward shock is not great, and it is offset somewhat by the slower collapse velocity of the outer surface of the thick shell, we can treat it approximately as a constant speed shock traversing the impacting shell. The converging shock can be treated by the classical model (see Section 3.7.3 Convergent Shocks). This allows us to estimate the minimum shell/levitated core mass ratio for efficient compression, the case in which the shock reaches the surface of the shell, and the center simultaneously.

If the shell and levitated core have identical densities and compressibilities, then the two shocks will have the same initial velocity (the velocity change behind the shock front in both cases will be exactly half the impact velocity). If the shell has thickness r_shell, then the shock will traverse the shell in time:

Eq. 4.1.6.2.3.2-1

$$t_{shell} = r_{shell}/v$$

If the levitated core has radius r_lcore, the shock will reach the center in time:

Eq. 4.1.6.2.3.2-2

$$t_{lcore} = (r_{lcore}/v)*alpha$$

Alpha is this case is the convergent shock scaling parameter (see Section 3.7.3). For a spherical implosion, and a gamma of 3 (approximately correct for most condensed matter, and for uranium and plutonium in particular), alpha is equal to 0.638 (the exact value will be somewhat higher than this).

Since we want t_shell = t_lcore:

Eq. 4.1.6.2.3.2-3

$$r_{shell} = alpha * r_{lcore} = 0.638\ r_{lcore}$$

That is, the thickness of the shell is smaller than the radius of the core by a factor of 0.638. But since volume is proportional to the cube of the radius:

Eq. 4.1.6.2.3.2-4

$$m_{shell} = density*(4*Pi/3)*[(r_{shell} + r_{lcore})^3 - r_{lcore}^3]$$

and

Eq. 4.1.6.2.3.2-5

$$m_{lcore} = density*(4*Pi/3)*r_{lcore}^3$$

This gives us the mass ratio:

Eq. 4.1.6.2.3.2-6

$$m_{shell}/m_{lcore} = ((1.638)^3 - 1^3)/1^3 = 3.4$$

Thus we want the impacting shell to have at least 3.4 times as much mass as the levitated core. The ratio used may be considerably larger.

Now it is important to realize that in principle the shell/levitated core mass ratio is unrelated to the tamper/fissile material mass ratio. The boundary between tamper and fissile material can be located in the shell (i.e. the shell is partly tamper and partly fissile, the levitated core entirely fissile), it can be located between the shell and core (i.e. the shell is tamper and the core is fissile), or it can be located in the core (i.e. the shell is tamper, and the core is partly tamper and partly fissile). The tamper/fissile material ratio is determined by neutron conservation, hydrodynamic confinement, and critical mass considerations.

It appears however that the initial practice of the US (starting with the Mk4 design and the Sandstone test series) was to design levitated core weapons so that the shell was the uranium tamper, and the levitated portion was a solid fissile core. The mass of the tamper would have been similar to that used in the Gadget (115 kg), a large enough mass to allow the use of different pit sizes and compositions while ensuring sufficient driver mass. These early pure fission bombs were designed to use a variety of pits to produce different yields, and to allow the composition (U-235/Pu-239 ratio) to be varied to match the actual production schedules of these materials.

Levitation is achieved by having some sort of support structure that will not disrupt the implosion symmetry. The most widely used approach seems to be the use of truncated hollow cones (or conically tapered thin walled tubes if you prefer), usually made out of aluminum. Six of these are

used, pairs on opposite sides of the levitated core for each axis of motion. Supporting wires (presumably under tension) have also been used.

The levitated core of the Hurricane device (the first British test) used "caltrops" (probably six of them) for support. A caltrops is a four pronged device originally used in the Middle Ages as an obstacle against soldiers and horses, and more recently against vehicle tires. Each of the prongs can be thought of as the vertex of a tetrahedron, with the point where they all join as the tetrahedral center. A caltrops has the property that no matter how you drop it, three of the prongs forms a tripod with the fourth prong pointed straight up. Dimples on the core might be used to seat the support prongs securely.

Another possibility is to use a strong light weight foam to fill the gap between shell and core (such foams have been produced at the Allied-Signal Kansas City Plant). A significant problem with using a foam support is that plastic foams are usually excellent thermal insulators, which could cause severe problems from self-heating in a plutonium levitated core.

A serious problem with hollow shell designs is the tensile stress generated by the Taylor wave (see Section 3.6.1.1.2 Free Surface Release Waves in Solids). As the release wave moves out from the inner shell surface, it encounters declining pressure due to the Taylor wave. The "velocity doubling" effect generates a pressure drop equal in magnitude to the shock peak pressure. If the pressure that the release wave encounters is below this pressure, a negative pressure (tension) is created (you can think of this as the faster moving part of the plate pulling the slower part along). This tensile stress builds up the farther back the release wave travels. If it exceeds the strength of the material it will fracture or "spall". This can cause the entire inner layer of material to peel off, or it may simply create a void. A new release wave will begin at the spall surface.

Spalling disrupts implosion symmetry and can also ruin the desired collision timing. It was primarily fears concerning spalling effects that prevented the use of levitated core designs in the first implosion bombs.

One approach to dealing with spalling is simply to make sure that excessive tensile stresses do not appear in the design. This requires strong materials, and at least one of the following:

1. relatively weak implosion pressures;
2. a highly suppressed Taylor wave; or
3. thin shells.

The first option is extremely undesirable, but the latter two can be used in viable designs.

Another approach is to adopt the "if you can't beat'em, join'em" strategy. Instead of trying to prevent separation in the shell, accommodation for the phenomenon is included in the design. This can be done by constructing the shell from separate layers. When the release wave reaches the boundary between shell layers (and tensile stress exists at that point), the inner layer will fly off the outer layer, and a new release wave will begin. This will create a series of imploding shells, separated by gaps.

As each shell layer converges toward the center, the inner surface will accelerate while the outer surface will decelerate. This will tend to bring the layers back together. If they do not rejoin before impact occurs with the core, a complicated arrangement of shocks may develop. The design possibilities for using these multiple shocks will not be considered here.

The concept of the levitated core and colliding shells can be extended to multiple levitation - having one collapsing shell collide with a second, which then collides with the levitated core. The outer shell, due to the concentration of momentum in its inner surface and the effects of elastic collision, could enhance the the velocity of the inner shell. This idea requires a large diameter system to be practical. It is possible that the "Type D" pit (that is, the hardware located between the explosive and fissile core) developed in the early fifties for the 60 inch diameter HE assemblies then in the US arsenal was such a system. It considerably increased explosive yields with identical cores.

It seems almost certain that the most efficient kiloton range pure fission bomb ever tested - the Hamlet device detonated in Upshot-Knothole Harry (19 May 1953) - used multiple levitation. It was described as being the first "hollow core" device, presumably the use of a fissile core that itself was an outer shell and an inner levitated core. A TX-13D bomb assembly (a 60 inch implosion system using a Type D pit) was used with the core. The yield was 37 kt.

4.1.6.2.3.3 Thin Shell (Flying Plate) Designs

Thin shell, or flying plate designs, take the hollow core idea to an extreme. In these designs a very thin, but relatively large diameter shell is driven inward by the implosion system. As with the regular hollow core design, a levitated core in the center is used.

The advantages of a flying plate design are: a greatly increased efficiency in the utilization of high explosive energy; and a higher collision speed - leading to faster insertion and greater compression for a given amount of explosive. Thin shell flying plate designs are standard now in the arsenals of the nuclear weapon states.

A thin plate, a few millimeters thick, is thinner than the Taylor wave of an explosive shock. The shock acceleration, followed by full release, is completed before the Taylor wave causes a significant pressure drop. The maximum initial shock acceleration is thus achieved.

Even greater energy transfer than this occurs however. When the release wave reaches the plate/explosive interface (completing the expansion and velocity doubling of the plate), a rarefaction wave propagates into the explosive gases. The gases expand, converting their internal energy into kinetic energy, and launching a new (but weaker) shock into the plate. A cyclic process thus develops in which a series of shocks of diminishing magnitude accelerate the plate to higher and higher velocities. If viewed from the inner surface, the observer would see a succession of velocity jumps of diminishing size and at lengthening intervals. The plate continues to accelerate over a distance of a few centimeters.

The maximum velocity achievable by this means can approach the escape velocity of the explosive gases, which is 8.5 km/sec for Comp B. Velocities up to 8 km/sec have been reported

using HMX-based explosives. This can be compared to the implosion velocity of the plutonium pit in the Gadget/Fat Man design, which was some 2 km/sec.

Optimum performance is found when a small gap (a few mm) separates the high explosive from the plate. Among other things, this gap reduces the strength of the Taylor wave. The gap may be an air space, but it is usually filled with a low impedance material (like a plastic).

The mass ratio between the explosive and the plate largely determines the system performance. For reasonable efficiency it is important to have a ratio r of at least 1 (HE mass/plate mass). At r=1 about 30% of the chemical energy in the explosive is transferred to the plate. Below r=1, the efficiency drops off rapidly. Efficiency reaches a maximum at r=2, when 35% of the energy is transferred.

Since a higher mass ratio means more energy available, the actual final velocity and energy in the plate increases monotonically with r, as shown in the table below. Higher values of r also cause the plate to approach its limiting value with somewhat shorter travel distances.

Table 4.1.6.2.3.3-1. Flying Plate Drive Efficiency

Plate/HE Mass Ratio (R)	Energy Fraction Transferred	Relative Velocity	Plate/Detonation Velocity Ratio
0.25	0.16	1.00	0.07
0.50	0.22	1.65	0.12
1.00	0.30	2.74	0.20
2.00	0.35	4.18	0.30
5.00	0.30	6.12	0.40
7.00	-	7.65	0.50
10.00	0.25	7.91	0.56

40.00	0.11	10.60	0.75

By the time the flying plate converges from a radius of 10-20 cm to collide with the levitated core, it is no longer a thin shell. The velocity difference that is inherent in thick shell collapse leads to a collision velocity of the inner surface that is higher than the average plate velocity. Collision velocities of experimental uranium systems of 8.5 km/sec have been reported.

The flying plate can be used in a variety of ways. It can be the collapsing shell of a levitated core design. Or it can be used as a driver which collides with, and transfers energy to a shell, which then implodes on to a levitated core.

4.1.6.2.3.4 Shock Buffers

Powerful shock waves can dissipate significant amounts of energy in entropic heating. Energy that contributes to entropy increase is lost to compression. This problem can be overcome by using a shock buffer.

A shock buffer is a layer of low impedance (i.e. low density) material that separates two denser layers. When a shock is driven into the buffer from one of the dense layers, a weaker shock of low pressure (but higher velocity) is created (see 3.6.1.1.3 Shock Waves at a Low Impedance Boundary). This shock is reflected at the opposite interface, driving a shock of increased pressure into the second dense layer. This shock is still weaker than the original shock however, and dissipates much less entropy.

A series of shock reflections ensue in the buffer, each one increases the pressure in the buffer, but by diminishing amounts (the pressure of the original shock is the limiting value). A series of shocks is driven into the second dense material, each successive shock creating a pressure jump of diminishing magnitude.

The shock buffer thus effectively splits the original powerful shock into a series of weaker ones, essentially eliminating entropic heating. The first two shocks produced account for most of the compression.

The following shocks tend to overtake the leading ones since they are travelling through compressed and accelerated material. Ideally, the shock sequence should be timed so that they all converge at the center of the system. The thickness of the buffer is selected so that this ideal is approached as closely as possible. The usual thickness is probably a few millimeters.

The buffer can be employed to cushion a plate collision also. In this case, the reflected shocks gradually decelerate the impactor (driver plate), and accelerate the driven plate, without dissipating heat. This converts a largely inelastic supersonic collision into an elastic one. If the mass of the driven plate is substantially lower than the mass of the driver, it can be accelerated to greater velocities than the original driver velocity. In principle an elastic collision can boost the

driven plate two as much as twice the velocity of the driver (if the driver/driven plate mass ratio is very large).

In practice this technique can transfer 65-80% of the driver energy to the driven plate, and provide driven plate velocities that are 50% greater than the driver velocity (or more). Since the explosive/plate mass ratio required for direct explosive drive increases very rapidly for velocities above 50% of the detonation velocity, the buffered plate collision method is the most efficient one for achieving velocities above this.

In an weapon implosion design a thin uranium or tungsten shell would probably be used as a driver.

Two likely low density materials for use as buffers are graphite and beryllium. Beryllium is an excellent neutron reflector which is commonly used in nuclear weapon designs for this reason. It thus may be a convenient shock buffer material that does double duty. Graphite is also a good neutron reflector. From information on manufacturing processes used at the Y-12 Plant at Oak Ridge, and the Allied-Signal Kansas City Plant, it is known that thin layers of graphite are used in the construction of nuclear weapons. The use of graphite as a shock buffer is a likely reason.

4.1.6.2.3.5 Cylindrical Implosion

The discussion of implosion has implicitly assumed a spherically symmetric implosion since this geometry is the simplest, and also the most efficient and widely used. Few changes are needed though to translate the discussion above to cylindrical geometry.

The changes required all relate to the differences in shock convergence in cylindrical geometry. There is a much lower degree of energy focusing during shock convergence, resulting in lower pressure increase for the same convergence ratio (reference radius/inner radius). A cylindrical solid core system would thus be much less effective in generating high pressures and compressions.

For a levitated core design, the shell/levitated core mass ratio must be recalculated. The appropriate value for alpha is 0.775 in this case, but the volume only increases by r^2, so:

Eq. 4.1.6.2.3.4-1

$$m_{shell}/m_{lcore} = ((1.775)^2 - 1^2)/1^2 = 2.15$$

The possibility of producing cylindrical implosion by methods that do not work for spherical geometries deserves some comment however. The flying plate line charge systems described above (4.1.6.2.2.4 Cylindrical and Planar Shock Techniques) for initiating a cylindrical implosion shocks in high explosives can be used to drive flying plates directly. Such a single-stage system would probably not be capable of generating as fast an implosion as a two stage system; one in which the first plate initiates a convergent detonation which then drives a second flying plate. A

single stage system would be simpler to develop and build, and potentially lighter and more compact however.

Cylindrical implosion systems are easier to develop that spherical ones. This largely because they are easier to observe. Axial access to the system is available during the implosion, allowing photographic and electronic observation and measurement. Cylindrical test systems were used to develop the implosion lens technology at Los Alamos that was later applied to the spherical bomb design.

4.1.6.2.3.6 Planar Implosion

Planar implosion superficially resembles the gun assembly method - one body is propelled toward another to achieve assembly. The physics of the assembly process is completely different however, with shock compression replacing physical insertion. The planar implosion process is some two orders of magnitude faster than gun assembly, and can be used with materials with high neutron background (i.e. plutonium).

By analogy with spherical and cylindrical implosion, the natural name for this technique might be "linear implosion". This name is used for a different approach discussed below in Hybrid Assembly Techniques.

Most of the comments made above about implosion still apply after a fashion, but some ideas, like the levitated core, have little significance in this geometry. Planar implosion is attractive where a cylindrical system with a severe radius constraint exists.

Shock wave lenses for planar implosion are much easier to develop than in other geometries. A plane wave lens is used by itself, not as part of a multi-lens system. It is much easier to observe and measure the flat shock front, than the curved shocks in convergent systems. Finally, flat shocks fronts are stable while convergent ones are not. Although they tend to bend back at the edges due to energy loss, plane shock fronts actually tend to flatten out by themselves if irregularities occur.

4.1.6.3 Hybrid Assembly Techniques

For special applications, assembly techniques that do not fit neatly in the previously discussed categories may be used.

4.1.6.3.1 Complex Guns

Additional improvements in gun system performance are possible by combining implosion with gun assembly. The implosion system here would be a very weak one - a layer of explosive to collapse a ring of fissile material or dense tamper on to the gun assembled core. This would allow further increases in the amount of fissile material used, and generate modest efficiency gains through small compression factors. A significant increase in insertion speed is also possible, which may be important where battlefield neutron sources may cause predetonation

(this may make the technique especially attractive for artillery shell use). Complex gun approaches have reportedly been used in Soviet artillery shell designs.

4.1.6.3.2 Linear Implosion

In weapons with severe size (especially radius) and mass constraints (like artillery shells) some technique other than gun assembly may be desired. For example, plutonium cannot be used in guns at all so a plutonium fueled artillery shell requires some other approach.

A low density, non-spherical, fissile mass can be squeezed and deformed into a supercritical configuration by high explosives without using neat, symmetric implosion designs. The technique of linear implosion, developed at LLNL, apparently accomplishes this by embedding an elliptical or football shaped mass in a cylinder of explosive, which is then initiated at each end. The detonation wave travels along the cylinder, deforming the fissile mass into a spherical form. Extensive experimentation is likely to be required to develop this into a usable technique.

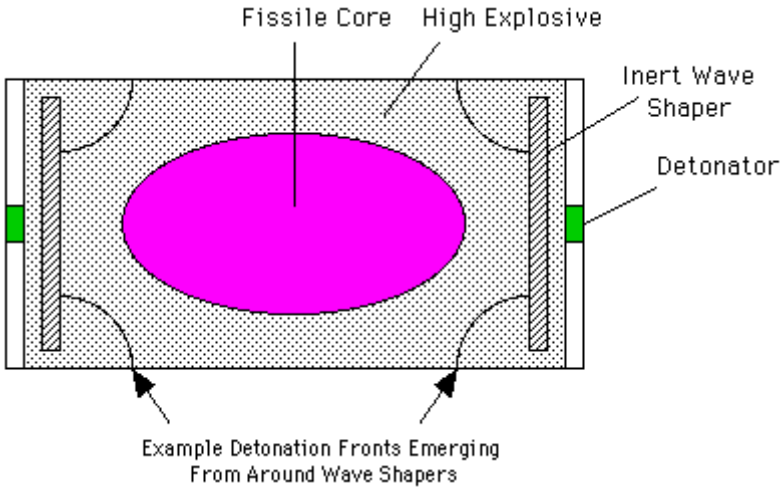

Three physical phenomenon may contribute to reactivity insertion:

- density increase due to collapsing voids in the core;
- density increase from phase transformations (if delta-phase plutonium is used); and
- reduction in surface area by deformation into a sphere (or approximate sphere).

Since the detonation generated pressure are transient, and affect different parts of the mass at different times, compression to greater than normal densities do not occur. The reactivity insertion then is likely to be rather small, and weapon efficiency quite low (which can be offset by boosting). The use of metastable delta-phase plutonium alloys is especially attractive in this type of design. A rather weak impulse is sufficient to irreversibly collapse it into the alpha phase, giving a density increase of 23%.

The supercritical mass formed by linear implosion is stable - it does not disassemble or expand once the implosion is completed. This relieves the requirement for a modulated neutron initiator, since spontaneous fission (or a calibrated continuous neutron source) can assure detonation. If desired, a low intensity initiator of the polonium/beryllium type can no doubt be used.

Special initiation patterns may be advantageous in this design, such as annual initiation - where the HE cylinder is initiated along the rim of each end to create a convergent shock wave propagating up the cylinder.

4.1.7 Nuclear Design Principles

The design of the nuclear systems of fission weapons naturally divides into several areas - fissionable materials, core compositions, reflectors, tampers, and neutron initiating techniques.

4.1.7.1 Fissile Materials

In the nuclear weapons community a distinction is made between "fissile" and "fissionable". Fissile means a material that can be induced to fission by neutrons of energy - fast or slow. These materials always have fairly high average cross sections for the fission spectrum neutrons of interest in fission explosive devices. Fissionable simply means that the material can be induced to fission by neutrons of a sufficiently high energy. As examples, U-235 is fissile, but U-238 is only fissionable.

There are three principal fissile isotopes available for designing nuclear explosives: U-235, Pu-239, and U-233. There are other fissile isotopes that can be used in principle, but various factors (like cost, or half-life, or critical mass size) that prevent them from being serious candidates. Of course none of the fissile isotopes mentioned above is actually available in pure form. All actual fissile materials are a mixture of various isotopes, the proportion of different isotopes can have important consequences in weapon design.

The discussion of these materials will be limited here to the key nuclear properties of isotope mixtures commonly available for use in weapons. The reader is advised to turn to Section 6 - Nuclear Materials for more lengthy and detailed discussions of isotopes, and material properties. See also Table 4.1.2-1 for comparative nuclear properties for the three isotopes.

4.1.7.1.1 Highly Enriched Uranium (HEU)

Highly enriched uranium (HEU) is produced by processing natural uranium with isotopic separation techniques. Natural uranium consists of 99.2836% U-238, 0.7110% U-235, and 0.0054% U-234 (by mass). Enrichment processes increase the proportion of light isotopes (U-235 and U-234) to heavy ones (U-238). Enriched uranium thus contains a higher percentage of U-235 (and U-234) than natural uranium, but all three isotopes are always present in significant concentrations. The term "HEU" usually refers to uranium with a U-235 of 20% or more. Uranium known to have been used in fission weapon designs ranges in enrichment from 80-93.5%. In the US uranium with enrichment around 93.5% is sometimes called Oralloy (abbreviated Oy) for historical reasons (Oralloy, or Oak Ridge ALLOY, was a WWII codename for weapons grade HEU). As much as half of the US weapon stockpile HEU has an enrichment in the range of 20-80%. This material is probably used in thermonuclear weapon designs.

The techniques which have actually been used for producing HEU are gaseous diffusion, gas centrifuges, electromagnetic enrichment (Calutrons), and aerodynamic (nozzle/vortex)

enrichment. Other enrichment processes have been used, some even as part of an overall enrichment system that produced weapons grade HEU, but none are suitable for the producing the highly enriched product. The original HEU production process used by the Manhattan Project relied on Calutrons, these were discontinued at the end of 1946. From that time on the dominant production process for HEU throughout the world has been gaseous diffusion. The vast majority of the HEU that has been produced to date, and nearly all that has been used in weapons, has been produced through gaseous diffusion. Although it is enormously more energy efficient, the only countries to have built or used HEU production facilities using gas centrifuges has been the Soviet Union, Pakistan, and The United Kingdom. Pakistan's production has been very small, the United Kingdom apparently has never operated there facility for HEU production.

High enrichment is important for reducing the required weapon critical mass, and for boosting the maximum alpha value for the material. The effect of enrichment on critical mass can be seen in the following table:

Figure 4.1.7.1.1. Uranium Critical Masses for Various Enrichments and Reflectors

total kg/U-235 content kg (density = 18.9)

Enrichment	Reflector		
(% U-235)	None	Nat. U	Be
		10 cm	10 cm
93.5	48.0/44.5	18.4/17.2	14.1/13.5
90.0	53.8/48.4	20.8/18.7	15.5/14.0
80.0	68. /54.4	26.5/21.2	19.3/15.4
70.0	86. /60.2	33. /23.1	24.1/16.9
60.0	120 /72.	45. /27.	32. /19.2
50.0	170 /85.	65. /33.	45. /23.
40.0	250 /100	100 /40.	70. /28.
30.0	440 /132	190 /57.	130 /39.
20.0	800 /160	370 /74	245 /49.

The total critical mass, and the critical mass of contained U-235 are both shown. The increase in critical mass with lower enrichment is of course less pronounced when calculated by U-235

content. Even with equivalent critical masses present, lower enrichment reduces yield per kg of U-235 by reducing the maximum alpha. This is due to the non-fission neutron capture cross section of U-238, and the softening of the neutron spectrum through inelastic scattering (see the discussion of U-238 as a neutron reflector below for more details about this).

U-238 has a spontaneous fission rate that is 35 times higher than U-235. It thus accounts for essentially all neutron emissions from even the most highly enriched HEU. The spontaneous fission rate in uranium (SF/kg-sec) of varying enrichment can be calculated by:

SF Rate = (fraction U-235)*0.16 + (1 - (fraction U-235))*5.5

For 93.5% HEU this rate (0.5 n/sec-kg) is low enough that large amounts can be used in weapon designs without concern for predetonation. If used in the Little Boy design (which actually used 80% enriched uranium, however) it would produce only one neutron every 31 milliseconds on average. No problem exists for any design up to the limiting size of gun-type weapons. 50% HEU on the other hand would be difficult to use in a gun-type weapon. A beryllium reflector would minimize the mass (and thus the amount of U-238 present), but to have a reasonable amount HEU present (e.g. 2.5 critical masses) would produce one neutron every 3.2 millisecs, making predetonation a significant prospect. The rate is never high enough though to make a significant difference for implosion assembly.

4.1.7.1.2 Plutonium

Plutonium is produced by neutron bombardment of U-238, which captures a neutron to form U-239. The U-239 then decays into neptunium-239, which decays in turn to form Pu-239. Since the vast majority of nuclear reactors use low enriched uranium fuel (< 20% U-235, 3-4% typically for commercial reactors), they also contain large amounts of U-238. Plutonium production is thus an inevitable consequence of operation in most reactors.

Pu-239 is the principal isotope produced, and is the most desired isotope for use in weapons or as a nuclear fuel. Multiple captures and other side reactions invariably produce an isotope mixture however. The principal contaminating isotope is always Pu-240, formed by non-fission neutron capture by Pu-239. The exposure of U-238 to neutron irradiation is measured by the fuel "burn-up", the number of megawatt-days (thermal) per tonne of fuel. The higher the burn-up, the greater the percentage of contaminating isotopes. Weapon production reactors use fuel burn-ups of 600-1000 MWD/tonne, light water power reactors have a typical design burn-up of 33000 MWD/tonne, and have been pushed to 45000 MWD/tonne by using higher enrichment fuel.

Plutonium is commonly divided into categories based on the Pu-240 content:

- < 3% Super grade
- 3-7% Weapons grade (normally 6-6.5%)
- 7-19% Fuel grade

- > 19% Reactor grade (spent fuel of LW power reactors)

The first US plutonium weapon (Fat Man) used plutonium with a Pu-240 content of only 0.9%, largely due to the hurried production schedule (only 100 MWD/tonne irradiations were used to get the plutonium out of the pile and into bombs quickly). Modern US nuclear weapons use weapons grade plutonium with a nominal 6.5% Pu-240 content. A lower Pu-240 content is not necessary for correct weapon functioning and increases the cost. The US has produced low-burnup supergrade plutonium to blend with higher burn-up feedstocks to produce weapons grade material. Plutonium produced in power reactors varies in composition, but its isotope profile remains broadly similar. If U-238 is exposed to extremely high burn-ups as in some fast breeder reactor designs (100,000 MWD/tonne), or if plutonium is separated from spent fuel and used as fuel in other reactors, it tends toward an equilibrium composition.

Representative plutonium compositions are:

	Pu-238	Pu-239	Pu-240	Pu-241	Pu-242
Weapon Grade	0.0%	93.6%	5.8%	0.6%	0.0%
	0.0%	92.8%	6.5%	0.7%	0.0%
Reactor grade	2.0%	61.0%	24.0%	10.0%	3.0%
Equilibrium	4.0%	32.0%	34.0%	15.0%	15.0%

These isotopes do not decay at the same rate, so the isotopic composition of plutonium changes with time (this is also true of HEU, but the decay process there is so slow as to be unimportant). The shortest lived isotopes found in weapon, fuel, or reactor grade plutonium in significant quantities are Pu-241 (13.2 yr) and Pu-238 (86.4 yr). The other isotopes have half-lives in the thousands of years and thus undergo little change over a human lifespan. The decay of Pu-241 (to americium-241) is of particular significance in weapons, since weapons grade plutonium contains no Pu-238 to speak of.

To understand the significance of these composition variations, we need to look at two principal factors: the critical mass size, and the spontaneous fission rate. An additional factor, decay self-heating, will be considered but is much less important.

Below are the estimated bare (unreflected) critical masses (kg) for spheres of pure plutonium isotopes in the alpha phase (and americium-241, since it is formed in weapons grade plutonium):

Pu-238	9 kg
Pu-239	10 kg

Pu-240 40 kg

Pu-241 12 kg

Pu-242 90 kg

Am-241 114 kg

The most striking thing about this table is that they all have critical masses! In contrast U-238 (or natural uranium, or even LEU) has no critical mass since it is incapable of supporting a fast fission chain reaction. This means that regardless of isotopic composition, plutonium will produce a nuclear explosion if it can be assembled into a supercritical mass fast enough.

Next observe that the critical masses for Pu-239 and Pu-241 are nearly the same, while the critical masses for Pu-240 and 242 are both several times higher. Because of this disparity, Pu-239 and Pu-241 tend to dominate the fissionability of any mixture, and it is commonplace in the literature to talk about these two isotopes as "fissile", while Pu-240 and 242 are termed "non-fissile". However it is not really true that 240 and 242 are non-fissile, which has an important consequence (shown in the table below):

Figure 4.1.7.1.2 Critical Masses for Plutonium of Various Compositions

total kg/Pu-239 content kg), density = 19.4

Isotopic Composition		Reflector	
atomic %		None	10 cm nat. U
239	240		
100%	0%	10.5/10.5	4.4/4.4
90%	10%	11.5/10.3	4.8/4.3
80%	20%	12.6/10.0	5.4/4.3
70%	30%	13.9/ 9.7	6.1/4.3
60%	40%	15.4/ 9.2	7.0/4.2
50%	50%	17.2/ 8.6	8.0/4.0
40%	60%	20.0/ 8.0	9.2/3.7
20%	80%	28.4/ 5.7	13. /2.6

0% 100% 40. / 0.0 20. /0.0

We can see that while the critical mass increases with declining "fissile" isotope content, the mass of Pu-239 present in each critical system diminishes. This is the exact opposite of the effect of isotopic dilution in uranium. In the range of isotopic compositions encountered in normal reactor produced plutonium, the content of Pu-239 in the reflected critical assemblies scarcely change at all. Thus regardless of isotopic composition, we can estimate the approximate critical mass based solely on the quantities of Pu-239, Pu-241 (and Pu-238) in the assembly.

Pu-242, having a higher critical mass, is a more effective diluent but it is only a minor constituent compared to Pu-240 in most isotopic mixtures. Even if Pu-242 is considered as the main diluent, the picture remains broadly similar.

The reason a relatively low concentration of Pu-240 is tolerable in weapon grade plutonium is due to the emission of neutrons through spontaneous fission. A high performance fission weapon is designed to initiate the fission reaction close to the maximum possible compression achievable by the implosion system, and predetonation must be avoided. The fastest achievable insertion rate is probably about 1 microsecond, it was 4.7 microseconds in Fat Man, and many designs will fall somewhere in the middle of this range.

We can calculate the spontaneous fission rate in a mass of plutonium with the following formula:

SF Rate (SF/kg-sec) = (%Pu-238)*1.3x10^4 + (%Pu-239)*1.01x10^-1 +

(%Pu-240)*4.52x10^3 + (%Pu-242)*8.1x10^3

For the 6.2 kg of plutonium (about 1% Pu-240) in Fat Man this is about 25,000 fissions/sec (or one every 40 microseconds). A weapon made with 4.5 kg of 6.5% Pu-240 weapon grade plutonium undergoes fission at a rate of 132,000 fission/sec (one every 7.6 microseconds). In an advanced design the window of vulnerability, in which a neutron injection will substantially reduce yield, might be as small as 0.5 microseconds, in this case weapon grade plutonium would produce only a 7% chance of substandard yield.

Even the plutonium found in the discharged fuel of light water power reactors can be used in weapons however. With a composition of 2% Pu-238, 61% Pu-239, 24% Pu-240, 10% Pu-241, and 3% Pu-242 we can calculate a fission rate of 159,000 fissions/kg-sec. If 6-7 kg were required in a design, then the average rate would be about 1 fission/microsecond. A fast insertion would have a significant chance of no predetonation at all, and would produce a substantial yield (a few kt) even in a worst case.

The US actually tested a nuclear device made from plutonium with a Pu-240 content of >19% in 1962. The yield was less than 20 kt. Although this was first made public in 1977, the exact amount of Pu-240, yield, and the date of the test are still classified.

Plutonium produces a substantial amount of heat from radioactive decay. This amounts to 2.4 W/kg in weapon grade plutonium, and 14.5 W/kg in reactor grade plutonium. This can make plutonium much warmer than the surrounding environment, and consideration of this heating effect must be taken into account in weapon design to ensure that deleterious temperatures aren't reached under any envisioned operating conditions. Thin shell designs are naturally resistant to these effects however, due to the large surface area of the thin plutonium shell. It can cause problems in levitated cores though, since the pit will have little thermal contact with surrounding materials.

Self heating can be calculated from the following formula:

Q (W/kg) = (%Pu-238)*5.67 + (%Pu-239)*0.019 + (%Pu-240)*0.07 + (%Pu-241)*0.034 + (%Pu-242)*0.0015 + (%Am-241)*1.06

The extremely weak decay energy of Pu-241 produces little heating considering the very short half-life, but Pu-241 decay does alter the isotopic and chemical composition substantially over a course of several years. Half of it decays over 13.2 years, giving rise to americium-241. This is a short half-life radioisotope with energetic decay. As Pu-241 is converted into americium significant increases in self-heating increases and radiotoxicity occur; a very slight (and probably insignificant) decline in reactivity also occurs.

Perhaps most important consequence of americium buildup is its effect on the alloy composition. Americium is one of the elements that can serve as an alloying agent to stabilize plutonium in the delta phase. Since alloying agents for this purpose are usually present to the extent of about 3% (atomic) in plutonium, a 0.6% addition of a new alloying agent (americium) is a significant composition change. This is not a serious problem with weapon grade plutonium, although it does have to be taken into account when selecting the alloy. In reactor grade plutonium the effect is quite pronounced since the decay of Pu-241 can add 10% americium to the alloy over a couple of decades. This would undoubtedly have important effects on alloy density and strength.

When refurbishing nuclear weapons it has been routine practice to extract americium from the plutonium and refabricate the pit. This is apparently not essential. The US is currently not refabricating weapon pits, and won't in significant numbers for several more years. Since weapon grade plutonium production has been shut down in the US, Russia, the UK, and France, the remaining supply of this material will become essentially free of Pu-241 (and Am-241 after reprocessing) over the next few decades.

4.1.7.1.2.1 Plutonium Oxide

Any sophisticated weapon design would use plutonium in the form of a metal, probably an alloy. The possibility of using plutonium (di)oxide (PuO_2) in a bomb design is of interest because the bulk of the separated plutonium existing worldwide is in this form. A terrorist group stealing plutonium from a repository might seek to use the oxide directly in a weapon.

Plutonium oxide is a bulky green powder as usually prepared. Its color may range from yellow to brown however. Oxygen has an extremely small neutron cross section, so plutonium oxide behaves essentially like a low density form of elemental plutonium. The maximum (crystal) density for plutonium oxide is 11.45, but the bulk powder is usually much less dense. A loose, unconsolidated powder might have a density of only 3-4. When compacted under pressure, substantially higher densities are achievable, perhaps 5-6 depending on pressure used. When compacted under very high pressure and sintered the oxide can reach densities of 9.7-10.0

The critical mass of reactor grade plutonium is about 13.9 kg (unreflected), or 6.1 kg (10 cm nat. U) at a density of 19.4. A powder compact with a density of 8 would thus have a critical mass that is (19.4/8)^2 time higher: 82 kg (unreflected) and 36 kg (reflected), not counting the weight of the oxygen (which adds another 14%). If compressed to crystal density these values drop to 40 kg and 17.5 kg.

4.1.7.1.3 U-233

Uranium-233 is the same chemical element as U-235, but its nuclear properties are more closely akin to plutonium. Like plutonium it is an artificial isotope that must be bred in a nuclear reactor. Its critical mass is lower than U-235, and its material alpha value is higher, both are close to those of Pu-239. Its half-life and bulk radioactivity are much closer to those of Pu-239 than U-235 also.

U-233 has been studied as a possible weapons material since the early days of the Manhattan Project. It is attractive in designs where small amounts of efficient material are desirable, but the spontaneous fission rate of plutonium is a liability, such as small, compact fission weapons with low performance (and thus light weight) assembly systems. It does not seem to have been used much, if at all, in actual weapons by the US. It has been employed in many US tests however, possibly indicating its use in deployed weapons.

The reason for this is the difficulty of manufacture. It must be made by costly irradiation in reactors, but unlike plutonium, its fertile isotope (thorium-232) is not naturally part of uranium fuel. To produce significant quantities of U-233, a special production reactor is required that burns concentrated fissile material for fuel - either plutonium or moderately to highly enriched uranium. This further increases cost and inconvenience, making it more expensive even than plutonium (which also has the advantage of a substantially lower critical mass). Significant resources have been devoted to U-233 production in the US however. In the fifties, up to three breeder reactors were loaded with thorium at Savannah River for U-233 production, and a pilot-scale "Thorex" separation plant was built.

U-233 has some advantages over plutonium, principally its lower neutron emission background. Like other odd numbered fissile isotopes U-233 does not readily undergo spontaneous fission, also important is the fact that the adjacent even numbered isotopes have relatively low fission rates as well. The principal isotopic contaminants for U-233 is U-232, which is produced by an n,2n reaction during breeding. U-232 has a spontaneous fission rate almost 1000 times lower than Pu-240, and is normally present at much lower concentrations.

If appropriate precautions are taken to use low Th-230 containing thorium, and an appropriate breeding blanket/reactor design is used, then weapons-grade U-233 can be produced with U-232 levels of around 5 parts per million (0.0005%). Above 50 ppm (0.005%) of U-232 is considered low grade.

Due to the short half-life of U-232 (68.9 years) the alpha particle emission of normal U-233 is quite high, perhaps 3-6 times higher than in weapons grade plutonium. This makes alpha->n reactions involving light element impurities in the U-233 a possible issue. Even with low grade U-233, and very low chemical purity uranium the emission levels are not comparable to emissions of Pu-240 in weapon grade plutonium, but they may be high enough to preclude using impure U-233 in a gun assembly weapon. If purity levels of 1 ppm or better are maintained for key light elements (achievable back in the 1940s, and certainly readily obtainable today), then any normal isotopic grade of U-233 can be used in gun designs as well.

Although the U-232 contaminant produces significant amount of self-heating (718 W/kg), it is presnt to small a concentration to have a significant effect. A bare critical mass of low grade U-233 (16 kg) would emit 5.06 watts, 11% of it due to U-232 heating.

Potentially a more serious problem is due to the decay chain of U-232. It leads to a series of short-lived isotopes, some of which put out powerful gamma emissions. These emissions increase over a period of a couple of years after the U-233 is refined due to the accumulation of the longest lived intermediary, Th-228. A 10 kg sphere of weapons grade U-233 (5 ppm U-232) could be expected to reach 11 millirem/hr at 1 meter after 1 month, 0.11 rem/hr after 1 year, and 0.20 rem/hr after 2 years. Glove-box handling of such components, as is typical of weapons assembly and disassembly work, would quickly create worker safety problems. An annual 5 rem exposure limit would be exceeded with less than 25 hours of assembly work if 2-year old U-233 were used. Even 1 month old material would require limiting assembly duties to less than 10 hours per week.

Typical critical mass values for U-233 (98.25%, density 18.6) are:

	Reflector		
	None	Nat. U	Be
	5.3 cm	10 cm	4.2 cm
Mass(kg)	16	7.6	5.7	7.6

Self heating can be calculated from the following formula:

Q (W/kg) = (%U-232)*7.18 + (%U-233)*0.0027 + (%U-234)*0.0018

4.1.7.2 Composite Cores

If more than one type of fissile material is available (e.g. U-235 and plutonium, or U-235 and U-233) an attractive design option is to combine them within a single core design. This eliminates the need for multiple weapon designs, can provide synergistic benefits from the properties of the two materials, and result in optimal use of the total weapon-grade fissile material inventory.

U-235 is produced by isotope enrichment and is generally much cheaper than the reactor-bred Pu-239 or U-233 (typically 3-5 times cheaper). The latter two materials have higher maximum alpha values, making them more efficient nuclear explosives, and lower critical masses. Plutonium has the undesirable property of having a high neutron emission rate (causing predetonation). U-233 has the undesirable property of having a high gamma emission rate (causing health concerns).

By combining U-235 with Pu-239, or U-235 with U-233, the efficiency of the U-235 is increased, and the required mass for the core is reduced compared to pure U-235. On the other hand, the neutron or gamma emission rates are reduced compared to pure plutonium or U-233 cores, and are significantly cheaper as well.

When a higher alpha material is used with a lower alpha material, the high alpha material is always placed in the center. Two reasons can be given for this. First, the greatest overall alpha for the core is achieved if the high alpha material (with the fastest neutron multiplication rate) is placed where the neutron flux is highest (i.e. in the center). Second, the neutron leakage from the core is determined by the radius of the core as measured in mean free paths. By concentrating the material with the shortest MFP in a small volume in the center, the "size" of the core in MFPs is maximized, and neutron leakage minimized.

Composite cores can be used in any type of implosion system (solid core, levitated core, etc.). The ratio of plutonium to HEU used has generally been dictated by the relative inventories or production rates of the two materials. These designs have largely dropped out of use in the US (and probably Soviet/Russian) arsenal as low weight thermonuclear weapon designs came to dominate the stockpile.

4.1.7.3 Tampers and Reflectors

Although the term "tamper" has long been used to refer to both the effects of hydrodynamic confinement, and neutron reflection, I am careful to distinguish between these effects. I use the term "tamper" to refer exclusively to the confinement of the expanding fissile mass. I use

"reflector" to describe the enhancement of neutron conservation through back-scattering into the fissile core. One material may perform both functions, but the physical phenomenon are unrelated, and the material properties responsible for the two effects are largely distinct. In some designs one or the other function may be mostly absent, and in other designs different materials may be used to provide most of each benefit.

Since the efficiency of a fission device is critically dependent on the rate of neutron multiplication, the effect of neutron conservation due to a reflector is generally more important than the inertial confinement effect of a tamper in maximizing device efficiency.

4.1.7.3.1 Tampers Tamping is provided by a layer adjacent to the fissile mass. This layer dramatically reduces the rate at which the heated core material can expand by limiting its velocity to that of a high pressure shock wave (a six-fold reduction compared to the rate at which it could expand into a vacuum).

Two physical properties are required to accomplish this: high mass density, and optical opacity to the thermal radiation emitted by core. High mass density requires a high atomic mass, and a high atomic density. Since high atomic mass is closely correlated to high atomic number, and high atomic number confers optical opacity to the soft X-ray spectrum of the hot core, the second requirement is automatically taken care of.

An additional tamping effect is obtained from the fact that a layer of tamper about one optical thickness (x-ray mean free path) deep becomes heated to temperatures comparable to the bomb core. The hydrodynamic expansion thus begins at the boundary of this layer, not the actual core/tamper boundary. This increases the distance the rarefaction wave must travel to cause significant disassembly.

To be effective, a tamper must be in direct contact with the fissile core surface. The thickness of the tamper need not be very large though. The shock travels outward at about the same speed as the rarefaction wave travelling inward. This means that if the tamper thickness is equal to the radius of the core, then by the time the shock reaches the surface of the tamper, all of the core will be expanding and no more tamping effect can be obtained. Since an implosion compressed bomb core is on the order of 3 cm (for Pu-239 or U-233), a tamper thickness of 3 cm is usually plenty.

In selecting a tamper, some consideration must be given to the phenomenon of Rayleigh-Taylor instability (see Section 3.8). During the period of inward flow following the passage of a convergent shock wave, instability can arise if the tamper is less dense than the fissile core. This is affected by the pressure gradient, length of time of implosion, implosion symmetry, the initial smoothness of the tamper/core interface, and the density difference.

The ideal tamper would the densest available material. The ten densest elements are (in descending order):

Osmium 22.57

Iridium 22.42

Platinum 21.45

Rhenium 21.02

Neptunium 20.02

Plutonium 19.84

Gold 19.3

Tungsten 19.3

Uranium 18.95

Tantalum 16.65

Although the precious metals osmium, iridium, platinum, or gold might seem to be too valuable to seriously consider blowing up, they are actually much cheaper than the fissile materials used in weapon construction. The cost of weapon-grade fissile material is inherently high. The US is currently buying surplus HEU from Russia for US$24/g, weapon grade plutonium is said to be valued 5 times higher. In the late 1940s U-235 cost $150/g in then-year dollars (worth several times current dollars)! If the precious metals actually had unique capabilities for enhancing the efficiency of fissile material, it might indeed be cost effective to employ them. No one is known to have actually used any of these materials as a fission tamper however.

Rhenium is much cheaper than the precious metals, and is a serious contender for a tamper material. Neptunium is a transuranic that is no cheaper than plutonium, and is actually a candidate fissile material itself. It is thus not qualified to be considered a tamper, nor is the costly and fissile plutonium. Gold would not be seriously considered as a tamper since tungsten has identical density but is much cheaper (it has been used as a fusion tamper however). Natural and depleted uranium (DU) has been widely used as a tamper due in large part to valuable nuclear properties (discussed below). The cheapness of DU (effectively free) certainly doesn't hurt.

Tungsten carbide (WC), with a maximum density of 15.63 (14.7 is more typical of fabricated pieces), is not an outstanding tamper material, but it is high enough to merit consideration as a combined tamper/reflector material since it is a very good reflector.

In comparison two other elements normally though of as being dense do not measure up: mercury (13.54), and lead (11.35). Lead has been used as a fusion tamper in radiation implosion designs though, either as the pure element or as a lead-bismuth alloy.

4.1.7.3.2 Reflectors

The usefulness of a material as a reflector is principally determined by its mean free path for scattering. The shorter this value, the better the reflector.

To see the importance of a short MFP, consider the typical geometry of a bomb - a spherical fissile core, with radius r_core, surrounded by a spherical reflector. The average distance from the center of the assembly at which an escaping neutron is first scattered is r_core + MFP. If the scattering MFP for a reflector is comparable to r_core, the reflector volume in which scattering occurs is much larger than the volume of the core. The direction of scattering is essentially random, so under these conditions a scattered neutron is unlikely to reenter the core. Most that eventually do reenter will have scattered several times, traversing a distance that is a multiple of the MFP value. Reducing the value of MFP will considerably reduce the volume in which scattering occurs, and thus increase the likelihood that a neutron will reenter, and reduce the average path it will traverse before doing so.

Since the neutron population in the core is increasing very fast, approximately doubling in the time it takes a neutron to traverse one MFP, the importance of an average reflected neutron to the chain reaction is greatly diluted by the "time absorption" effect. It represents an older and thus less numerous neutron generation, which has been overwhelmed by more recent generations. This effect can be represented mathematically by including in the reflector a fictitious absorber whose absorption cross section is inversely proportional to the neutron velocity. Due to time absorption, as well as the effects of geometry, effectiveness of a reflector thus drops very rapidly with increasing MFP.

For a constant MFP, increasing reflector thickness also has a point of diminishing returns. Most of the benefit in critical mass reduction occurs with a reflector thickness of one 1 MFP. With 2 MFPs of reflector, the critical mass has usually dropped to within a few percent of its value for an infinitely thick reflector. Time absorption also causes the benefits of a reflector to drop off rapidly with thicknesses exceeding about one MFP. A very thick reflector offers few benefits over a relatively thin one.

Experimental data showing the variation of critical mass with reflector thickness can be misleading for evaluating reflector performance in weapons since critical systems are non-multiplying (alpha = 0). These experiments are useful when the reflector is relatively thin (a few centimeters), but thick reflector data is not meaningful. For example, consider the following critical mass data for beryllium reflected plutonium:

Table 4.1.7.3.2-1. Beryllium-Plutonium Reflector Savings

Beryllium Thickness (cm)	Alpha Phase Pu Critical Mass (d = 19.25) (kg)
0.00	10.47
5.22	5.43

8.17	4.66
13.0	3.93
21.0	3.22
32.0	2.47

The very low critical mass with a 32 cm reflector is meaningless in a high alpha system, it would behave instead as if the reflector were much thinner (and critical mass correspondingly higher). Little or no benefit is gained for reflectors thicker than 10 cm. Even a 10 cm reflector may offer slight advantage over one substantially thinner.

> [Note: The table above, combined with the 2 MFP rule for reflector effectiveness, might lead one to conclude that beryllium's MFP must be in the order of 16 cm. This is not true. Much of the benefit of very thick beryllium reflectors is due to its properties as a moderator, slowing down neutrons so that they are more effective in causing fission. This moderation effect is useless in a bomb since the effects of time absorption are severe for moderated neutrons.]

In the Fat Man bomb, the U-238 reflector was 7 cm thick since a thicker one would have been of no value. In assemblies with a low alpha, additional reflectivity benefits are seen with uranium reflectors exceeding 10 cm thick. To reduce the neutron travel time it is also important for the neutron reflector to be in close proximity to the fissile core, preferably in direct contact with it.

Since MFP decreases when the reflector is compressed, it is very beneficial to compress the reflector along with the fissile core.

Many elements have similar scattering microscopic cross sections for fission spectrum neutrons (2.5 - 3.5 barns). Consequently the MFP tends to correlate with atomic density. Some materials (uranium and tungsten for example) have unusually high scattering cross sections that compensate for a low atomic density.

The parameter c (the average number of secondaries per collision) is also significant. This is the same c mentioned earlier in connection with the alpha of fissile materials. In reflector materials the effective value of c over the spectrum of neutrons present is always less than 1. Only two reflector materials produce significant neutron multiplication: U-238 (from fast fission) and beryllium (from the Be-9 + n -> 2n + Be-8 reaction). Neutron multiplication in U-238 becomes significant when the neutron energy is above 1.5 MeV (about 40% of all fission neutrons), but a neutron energy of 4 MeV is necessary in beryllium. Further, U-238 produces more neutrons per reaction on average (2.5 vs 2). For fission spectrum neutrons this gives U-238 a value of c = 1.05, and Be a value of c = 1.03. Remember, this if for *fission spectrum neutrons*, i.e. neutron undergoing their first collision! The *effective* value is lower though since after one or more collisions the energy spectrum changes.

Each uranium fast fission neutron is considerably more significant in augmenting the chain reaction in the core, compared to beryllium multiplied neutrons, due to the higher energy of fast fission neutrons. U-238 fast fission is an energy producing reaction, and generates neutrons with an average energy of 2 MeV. The beryllium multiplication reaction *absorbs* energy (1.665 MeV per reaction) and thus produces slow, low energy neutrons for whom time absorption is especially severe. The energy produced by U-238 fast fission can also significantly augment the yield of a fission bomb. It is estimated that 20% of the yield of the Gadget/Fat Man design came from fast fission of the natural uranium tamper.

Both beryllium and uranium have negative characteristics in that they tend to reduce the energy of scattered neutrons (and reduce the effective value of c below 1). In beryllium this is due to moderation - the transfer of energy from the neutron to an atomic nucleus through elastic scattering. In uranium it is due to inelastic scattering.

4.1.7.3.2.1 Moderation and Inelastic Scattering

The energy loss with moderation is a proportional one - each collision robs the neutron of the same average fraction of its remaining energy. This fraction is determined by the atomic weight of the nucleus:

E_collision/E_initial = Exp(-epsilon)

the constant epsilon being calculated from:

epsilon = 1 + ((A - 1)^2 * ln((A - 1)/(A + 1))/(2*A))

where A is the atomic number. The equation is undefined when A=1, but taking the limit as it approaches 1 gives the value for light hydrogen which is epsilon=1. If A is larger than 5 or so then it can be approximated by:

epsilon ~= 2/(A + 2/3).

Epsilon values for some light isotopes of interest are:

A	Isotopes	Epsilon
1	H	1.000
2	D	0.725
3	T, He-3	0.538

4 He-4 0.425

6 Li-6 0.299

7 Li-7 0.260

9 Be-9 0.207

10 B-10 0.187

12 C-12 0.158

Since epsilon is close to zero when A is large, we can easily see that moderation is significant only for light atoms. The atomic weight of beryllium (9) is light enough to make this effect significant.

The average number of collisions n required to reduce a neutron of energy

E_initial to E_final can be expressed by:

n = (1/epsilon) * ln(E_initial/E_final)

Since A=9 for beryllium, it takes 3.35 collisions to reduce neutron energy by half. The average number of collisions for a neutron reentering the fissile mass will likely be substantially higher than this, unless the reflector is thin (in which case most of the neutrons will escape without reflection). For comparison carbon (A=12) takes 4.39 collisions to achieve similar moderation, iron (A=56) takes 19.6, and U-238 takes 165.

Clearly heavy atoms do not cause significant moderation. However they can experience another phenomenon called inelastic scattering that also absorbs energy from neutrons. In inelastic scattering, the collision excites the nucleus into a higher energy state, stealing the energy from the neutron. The excited nucleus quickly drops back to its ground state, producing an x-ray. Inelastic scattering is mostly important only in very heavy nuclei that have many excitation states (like tungsten and uranium). The effect drops off rapidly with atomic mass.

In balance, the energy loss by moderation in beryllium is more serious than the energy loss by inelastic scattering in uranium. This is partly due to the fact that every elastic collision reduces neutron energy, while only some collisions produce inelastic scattering.

4.1.7.3.2.2 Comparison of Reflector Materials

Below is a list of candidate materials, and their atomic densities. The list includes the six highest atomic density pure elements (C - in two allotropic forms, Be, Ni, Co, Fe, and Cu), and a

number of compounds that are notable for having high atomic densities. Atomic densities for the major tampers materials are also shown.

Table 4.1.7.3.2.2-1. Candidate Reflector Materials

Cross sections and MFPs are for fission spectrum neutrons

Reflector Material	At. Density moles/cm^3	Avg. Cross. barns	MFP cm
Carbon (C,diamond)	0.292	2.37	2.40
Beryllium Oxide (BeO)	0.241	2.79	2.47
Beryllium (Be)	0.205	2.83	2.86
Beryllium Carbide (BeC)	0.190	2.60	3.36
Carbon (C, graphite)	0.188	2.37	3.73
Water (H2O)	0.167	3.54	2.81
Nickel (Ni)	0.152	3.84	2.85
Tungsten Carbide (WC)	0.150	4.55	2.43
Cobalt (Co)	0.148	3.68	3.05
Iron (Fe)	0.141	3.66	3.22
Copper (Cu)	0.141	3.65	3.23
...			
Osmium (Os)	0.118		
Iridium (Ir)	0.117		
Rhenium (Re)	0.110		
Platinum (Pt)	0.110		
Tungsten (W)	0.105	6.73	2.35
Gold (Au)	0.098		
Plutonium (Pu)	0.083		

Uranium (U)	0.080	7.79	2.66
Mercury (Hg)	0.068		
Lead (Pb)	0.055		

From this list it can be seen that the highest atomic density materials consist of light elements. Some compounds achieve higher atomic densities than pure elements by packing together atoms of different sizes. Thus BeO is denser (in both mass and moles/cm^3) than Be, and WC is denser than W (only in moles/cm^3).

Using critical mass data, some of these materials can be ordered by reflector efficiency. In the ordering below X > Y means X is a better reflector than Y, and (X > Y) means that though X is better than Y, the difference is so slight that they are nearly equal (MFPs are shown below each material):

Be > (BeO > WC) > U > W > Cu > H2O > (Graphite > Fe)

2.86 2.47 2.28 2.66 2.43 3.23 2.82 3.73 3.22

From this the general trend of lower MFPs for better reflectors is visible, but is not extremely strong. The effects of neutron multiplication and moderation are largely responsible. As noted earlier this ranking, made using critical assemblies, tends to overvalue beryllium somewhat with respect to use in weapons. Nonetheless beryllium is still by and large the best reflector, especially when low mass is desirable. Uranium and tungsten carbide are the best compromise reflector/tampers.

Carbon is a fairly good neutron reflector. It has the disadvantage of being a light element that moderates neutrons, but being heavier than beryllium (At Wt 12 vs 9) it moderates somewhat less. When used as a shock buffer, additional significant benefits from neutron reflection can be obtained. The singularly high atomic density and short MFP for diamond makes it an interesting material. Before dismissing the possibility out of hand as ridiculous, given its cost, it should be noted that synthetic industrial diamond cost only $2500/kg, far less than the fissile material used in the core. It can also be formed into high density compacts.

Iron is a surprisingly good reflector, though not good enough to be considered for this use in sophisticated designs. It may be important due to its use as a structural material - as in the casing of a nuclear artillery shell, or the barrel of gun-type weapon.

With a 4.6 cm radius core the following reflector thicknesses have been found to be equally effective:

Be 4.2 cm

U 5.3 cm

W 5.8 cm

Graphite 10. Cm

Viewed from the other perspective (variation in critical mass with identical thicknesses of different materials) we get:

Table 4.1.7.3.2.2-2. Critical Mass for 93.5% U-235 (kg)

Material	Reflector Thicknesses		
	2.54 cm	5.08 cm	10.16 cm
Be	29.2	20.8	14.1
BeO	-	21.3	15.5
WC	-	21.3	16.5
U	30.8	23.5	18.4
W	31.2	24.1	19.4
H2O		24	22.9
Cu	32.4	25.4	20.7
Graphite	35.5	29.5	24.2
Fe	36.0	29.5	25.3

Below is a plot showing the change of Oralloy critical mass with reflector thickness graphically (taken from LA-10860-MS, Critical Dimensions of Systems Containing 235-U, 239-Pu, and 233-U; 1986 Rev.):

Figure 4.1.7.3.2.2-2. Critical Mass for Oralloy for Various Reflectors

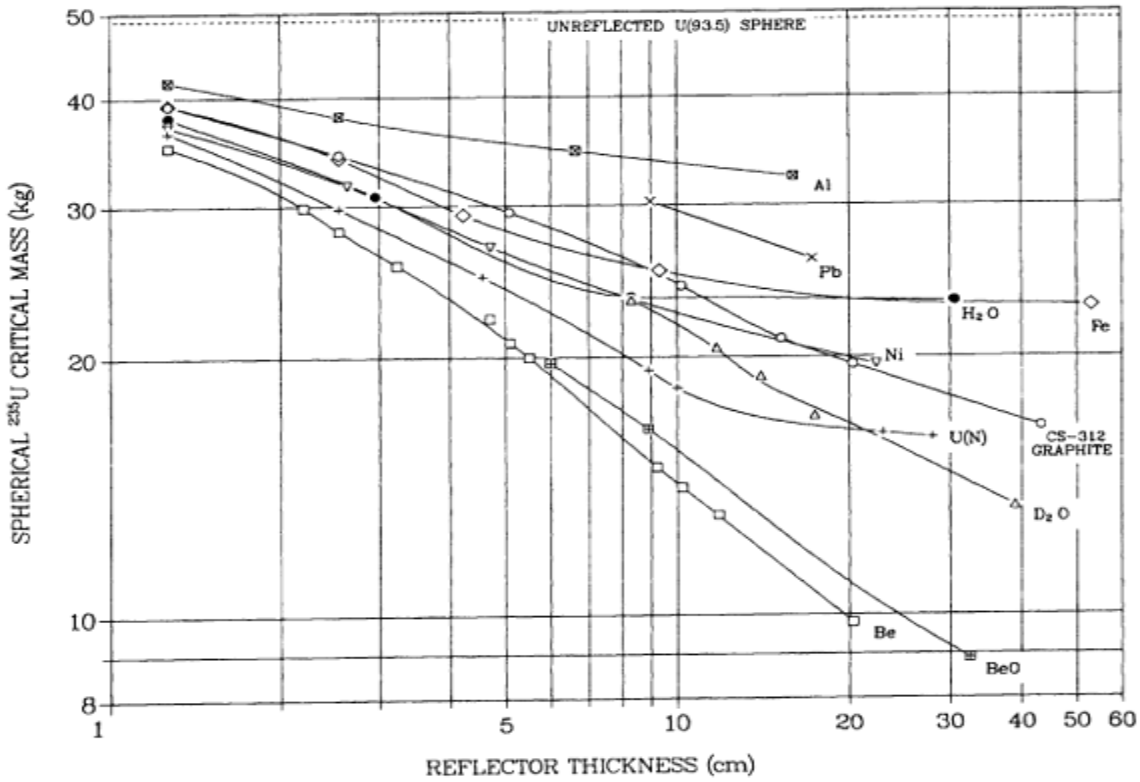

The variation of plutonium and U-233 critical masses with reflector thickness can be determined using the chart below (also taken from LA-10860-MS) with the above chart for Oralloy:

Figure 4.1.7.3.2.2-3. Plutonium/Oralloy and U-233/Oralloy Critical Mass Ratios for Various Reflectors

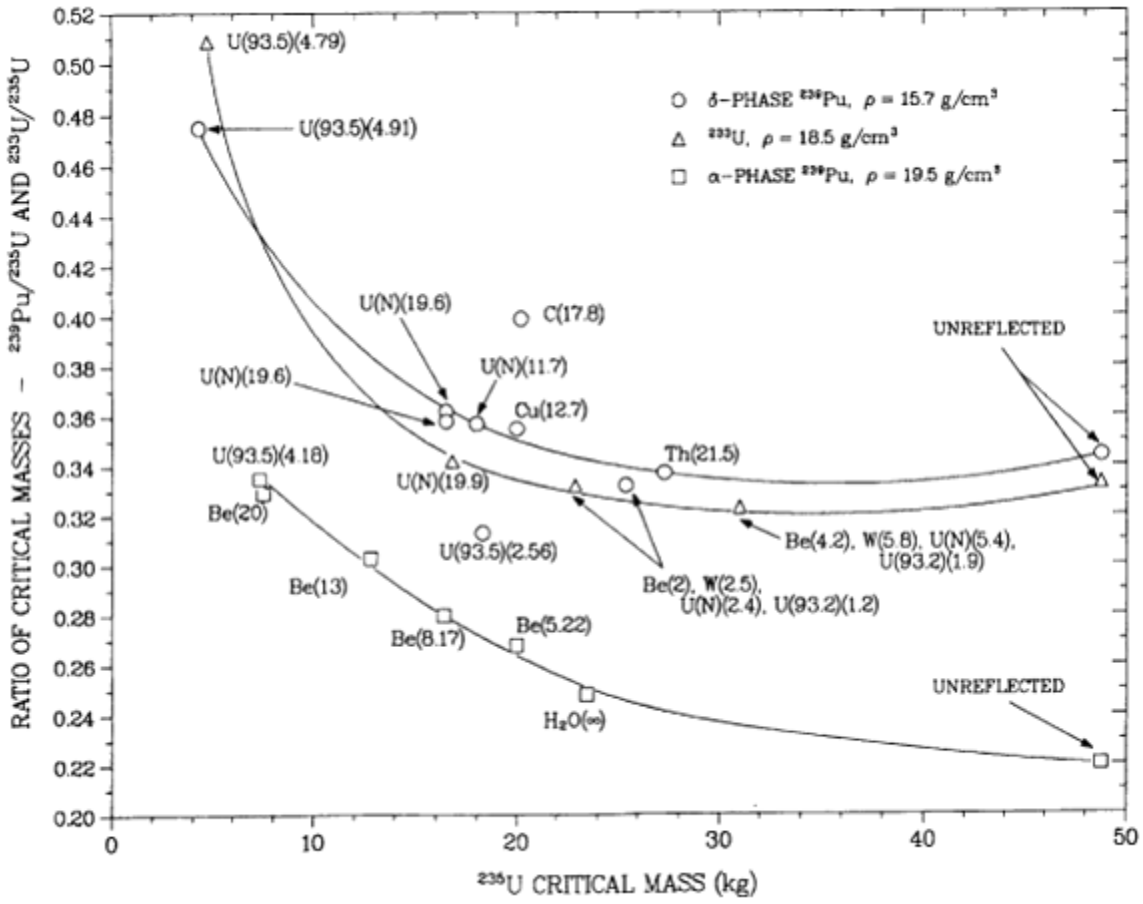

The variation of critical mass with reflector thickness is sometimes also expressed in terms of reflector savings, the reduction in critical radius for a given reflector thickness:

Table 4.1.7.3.2.2-3. Reflector Savings (cm) for Various Reflector/Fissile Material Combinations

	Fissile Material							
	93.5% U-235				Plutonium			
Reflector	Reflector Thicknesses (cm)				Reflector Thicknesses (cm)			
Material	1.27	2.54	5.08	10.16	1.27	2.54	5.08	10.16
Be	0.90	1.46	2.14	2.94	0.73	1.11	1.51	1.97
U	0.81	1.31	1.87	2.40	0.66	1.01	1.36	1.66
W	0.82	1.29	1.82	2.29	0.67	1.00	1.33	1.59
Fe	0.59	0.92	1.36	1.70	0.50	0.74	1.04	1.25

4.1.7.3.3 Combined Tamper/Reflector Systems

In most weapon designs, both the benefits of tamping and neutron reflection are desired. Two design options are available:

1. Use a compromise material that performs both functions; or
2. Use a layered system - an inner tamper and an outer reflector.

Designs for relatively heavy implosion bombs typically use U-238 (as natural or depleted uranium) as a compromise material. It is very good to excellent in both respects, and boosts yield as well. The Gadget/Fat Man design used a 120 kg natural uranium tamper (7 cm thick). All of the early U.S. implosion designs used uranium as a tamper/reflector. The spontaneous fission rate in U-238 precludes its use in gun-type designs.

The Little Boy weapon used tungsten carbide as a compromise material. Its density is fairly high, and it is an excellent neutron reflector (second only to beryllium among practical reflector materials). It is less dense than the uranium core, but since the Little Boy core was not compressed, Rayleigh-Taylor instability was not a factor in design. Tungsten metal was used in the South African gun-type weapons, this choice places greater emphasis on tamping over reflection, compared to tungsten carbide. It is interesting to note the dual-use restrictions placed on tungsten alloys and carbide:

> Parts made of tungsten, tungsten carbide, or tungsten alloys (>90% tungsten) having a mass >20 kg and a hollow cylindrical symmetry (including cylinder segments) with an inside diameter greater than 10cm but less than 30 cm.

This is clearly based on its use as a reflector in gun-type weapons.

Beryllium is used as a reflector in modern light weight fission warheads, and thermonuclear triggers. It has special value for triggers since it is essentially transparent to thermal radiation emitted by the core. It is a very efficient reflector for its mass, the best available. But due to its extremely low mass density, it is nearly useless as a tamper. In boosted designs tamping may be unnecessary, but it is also possible to insert a (thin) tamper layer between the core and beryllium reflector). The n,2n reaction is also useful in boosted designs, since that fraction of fusion neutrons that escape the core without capture or substantial scatter still retain enough energy to release reasonably energetic neutrons in the reflector. Beryllium has relatively high compressibility, which may also add to its effectiveness as a reflector.

It is also interesting to note that the Allied-Signal Kansas City Plant has developed a capability for depositing tungsten-rhenium films up to 4 mm thick. This would be a nearly ideal material and thickness for a tamper in a beryllium reflected flying plate implosion design. By alloying rhenium with tungsten, the density of the tungsten can be increased (so that it matches or exceeds the density of alpha phase plutonium), and the ductility and workability of tungsten is improved. Notable confirmation of this comes form the 31 kt Schooner cratering test in 1968

(part of the Plowshare program). Some of the most prominent radionuclides in the debris cloud were radioactive isotopes of tungsten and rhenium.

It is also possible that uranium foils known to have been manufactured for weapons were used as tampers in flying plate designs.

4.1.8 Fission Initiation Techniques

Once a supercritical mass is assembled, neutrons must be injected to start the chain reaction.

This is not really a problem for a gun type weapon, since the design allows the supercritical mass to remain in the fully assembled state indefinitely. Eventually a neutron from the prevailing background is certain to cause a full yield explosion.

It is a major problem in an implosion bomb since the interval during which the bomb is near optimum criticality is quite short - both in absolute length (less than a microsecond), and also as a proportion of the time the bomb is in a critical state.

The first technique to be seriously considered for use in a weapon was simply to include a continuous neutron emitter, either a material with a high spontaneous fission rate, or an alpha emitter that knocks neutrons loose from beryllium mixed with it. Such an emitter produces neutrons randomly, but with a specific average rate. This inevitably creates a random distribution in initiation time and yield (called stochastic initiation). By tuning the average emission rate a balance between pre and post detonation can be achieved so that a high probability of a reasonably powerful (but uncertain) yield can be achieved. This idea was proposed for the Fat Man bomb at an early stage of development.

A far superior idea is to use a modulated neutron initiator - a neutron emitter or neutron generator that can be turned on at a specific time. This is a much more difficult approach to develop, regardless of the technique used. Modulated initiators can be either internal designs, which are placed inside the fissile pit and activated by the implosion wave, or external designs which are placed outside the fission assembly.

It should be noted that it is very desirable for an initiator to emit at least several neutrons during the optimum period, since a single neutron may be captured without causing fission. If a large number can be generated then the total length of the chain reaction can be significantly shortened. A pulse of 1 million neutrons could cut the total reaction length by 25% or so (approx. 100 nanoseconds), which may be useful for ensuring optimal efficiency.

4.1.8.1 Modulated Beryllium/Polonium Initiators

This general type of initiator was used in all of the early bomb designs. The fundamental idea is to trigger the generation of neutrons at the selected moment by mixing a strong alpha emitter with the element beryllium. About 1 time out of 30 million, when an alpha particle collides with a beryllium atom a neutron is knocked loose.

The key difficulty here is keeping the alpha emitter out of contact with the beryllium, and then achieving sufficiently rapid and complete mixing that a precisely timed burst of neutrons is emitted.

The very short range of alpha particles in solid matter (a few tens of microns) would make the first requirement relatively easy to achieve, except for one thing. Most strong alpha emitters also emit gamma rays, which penetrate many centimeters of solid matter and also occasionally knock loose neutrons. Finding a radioisotope with sufficiently low gamma emissions greatly restricts the range of choices. A suitable radioisotope must also have a relatively short half-life (no more than a few decades) so sufficient activity can be provided by a small amount, and be reasonably economical to produce.

One isotope appears to be the clear favorite when all these factors are considered: polonium 210. Although other alpha emitters have been considered, all radioisotope based modulated initiators appear to have used Po-210 as the alpha source. This isotope has a half-life of only 138.39 days though. On the one hand, this means a strong emitter alpha source can be quite small (50 curies, which emits 1.85 x 10^12 alphas/sec, weighs only 11 mg). On the other, the Po-210 disappears quickly and must be constantly replenished to maintain a standing arsenal. Polonium-208 and actinium-227 have also been considered for this role.

The second requirement: carefully timed, fast, efficient mixing, needs very clever designs for implosion weapons. After considering several proposals, a neutron initiator called "Urchin" or "screwball" was selected by Los Alamos for Gadget/Fat Man. All of the designs considered were based on placing the initiator at the center of the fissile mass, and using the arrival of the convergent shock to drive the mixing process. This insured that the entire mass was highly compressed (although perhaps not optimally compressed), and placed the initiator where the neutrons emitted would be most effective.

The Urchin was a sphere consisting of a hollow beryllium shell, with a solid spherical beryllium pellet nested inside. The polonium was deposited in layer between the shell and the pellet. Both the shell and the pellet were coated with a thin metal film to prevent the polonium (or its alpha particles) from reaching the beryllium. The mixing was brought about by using the Munroe Effect (also called the shaped charge, or hollow charge, effect): shock waves collide, powerful high velocity jets are formed. This effect was created by cutting parallel wedge-shaped groves in the inner surface of the shell. When the implosion shock collapsed these grooves, sheet-like beryllium jets would erupt through the polonium layer, and cause violent turbulence that would quickly mix the polonium and beryllium together.

By placing the small mass of polonium as a layer trapped between two relatively large masses of beryllium, the Urchin designers were hedging their bets. Even if the Monroe effect did not work as advertised, any mixing process or turbulence present would likely disrupt the carefully isolated polonium layer and cause it to mix.

The whole initiator weighed about 7 grams. The outer shell was 2 cm wide and 0.6 cm thick, the solid inner sphere was 0.8 cm wide. 15 parallel wedge-shaped grooves, each 2.09 mm deep,

were cut into the inner surface of the shell. Both the shell and the inner pellet were formed by hot pressing in a nickel carbonyl atmosphere, which deposited a nickel layer on the surfaces. The surfaces of the shell and central sphere were also coated with 0.1 mm of gold. Combined with the nickel layer, the gold film provided a barrier between the polonium and the beryllium.

50 curies polonium-210 (11 mg) was deposited on the grooves inside the shell and on the central sphere. This much polonium produces a thermal output of 0.1 watts, causing very noticeable warming in such a small object. Post war studies showed that no more than 10 curies still provided an acceptable initiation effect, allowing the manufacture of initiators that remained usable for up to a year.

Other designs for generating mixing have been considered. One design considered during or shortly after WWII used a spherical shell whose interior surface was covered with conical indentations. The shell was coated with a metal film, and polonium was deposited on the interior surface as in the Urchin design. In this design the cavity inside the hollow shell was empty, there was no central pellet. The principal advantage here is that the initiator could be made smaller while still being reliable. A shortcoming of the Urchin was that the Munroe effect is less robust in linear geometry. The formation of a jet when a wedge collapses depends on the apex angle and other factors, and could conceivably fail (its use may have been due to the more thorough study given the linear geometry by Fuchs during the war). The jet effect is quite robust in conical geometry however, the collapse of the conical pits producing high velocity jets of beryllium metal squirting into the cavity under nearly all conditions. Pyrimidal pits provide similar advantages, and have been used in hollow and central sphere equipped initiators.

The smaller TOM initiator (about 1 cm) that replaced the Urchin was probably based on the hollow conical pit (or tetrahedral pit) design. This design was proposed for use in 1948, but not put into production until January 1950 by Los Alamos. It was first tested (in a weapon test) in May 1951. One advantage of the TOM initiator was more efficient use of the polonium (more neutrons per gram of Po-210).

One sophisticated design that was developed and patented by Klaus Fuchs and Rubby Sherr during the Manhattan project was based on using the outgoing implosion rebound, rather than the incoming converging shock to accomplish mixing. This slight delay in initiation thus achieved was expected to allow significantly more compression to occur.

If internal initiators are used in fusion-boosted designs it is essential that they be quite small, the smaller the better (external initiation is best).

In gun-type weapons initiators are not strictly required, but may be desirable if the detonation time of the weapon needs to be precisely controlled. A low intensity polonium source can be used in this case, as can a simple system to bring the source and beryllium into contact upon impact by the bullet (like driving a beryllium foil coated piston into a sleeve coated with polonium).

4.1.8.2 External Neutron Initiators (ENIs)

These devices (sometimes called "neutron generators") rely on a miniature linear particle accelerator called a "pulse neutron tube" which collides deuterium and tritium nuclei together to generate high energy neutrons through a fusion reaction. The tube is an evacuated tube a few centimeters long with an ion source at one end, and an ion target at the other. The target contains one of the hydrogen isotopes adsorbed on its surface as a metal hydride (which isotope it is varies with the design).

When a current surge is applied to the ion source, an electrical arc creates a dense plasma of hydrogen isotope ions. This cloud of ions is then extracted from the source, and accelerated to an energy of 100-170 KeV by the potential gradient created by a high voltage acceleration electrode. Slamming into the target, a certain percentage of them fuse to release a burst of 14.1 MeV neutrons. These neutrons do not form a beam, they are emitted isotropically.

Early pulse neutron tubes used titanium hydride targets, but superior performance is obtained by using scandium hydride which is standard in current designs.

A representative tube design is the unclassified Milli-Second Pulse (MSP) tube developed at Sandia. It has a scandium tritide target, containing 7 curies of tritium as 5.85 mg of ScT2 deposited on a 9.9 cm^2 molybdenum backing. A 0.19-0.25 amp deuteron beam current produces about 4-5 x 10^7 neutrons/amp-microsecond in a 1.2 millisecond pulse with accelerator voltages of 130-150 KeV for a total of 1.2 x 10^10 neutrons per pulse. For comparison the classified Sandia model TC-655, which was developed for nuclear weapons, produced a nominal 3 x 10^9 neutron pulse.

A variety of ion source designs can be used. The MSP tube used a high current arc between a scandium deuteride cathode and an anode to vaporize and ionize deuterium. Other designs (like the duoplasmotron) may use an arc to ionize a hydrogen gas feed. The ion output current limits the intensity of the neutron pulse. Public domain ion source designs typically have a ion current limit of several amps. If we assume that the TC-655 achieved a 10 amp current from its ion source (the design of which is classified) then we can estimate an emission rate of up to 5 x 10^8 neutrons/microsecond in a pulse 6 microseconds long.

It is misleading though to think of a neutron tube as producing all its neutrons in a sudden burst. From the perspective of the fission process in a bomb core, it is not sudden at all. A typical core alpha is 100-400/microsecond, with corresponding neutron multiplication intervals of 2.5-10 nanoseconds. Any neutrons that enter the core in one multiplication interval will increase by a factor of e (2.7...) in the next, overwhelming the external neutron flux. From this point on, the fission process will proceed on its course unaffected by the ENI. Only the neutrons that enter the core during a single multiplication interval really count, and they count only insofar as they determine the time that the exponential chain reaction begins. Clearly, the vast majority of the neutrons in a 6 microsecond pulse are utterly irrelevant. The important factor in determining how effective an ENI is in precisely controlling the start of the chain reaction is the beam current intensity and how sharply and precisely it can be turned on. These are the design parameters that should be optimized in a weapon tube.

Note that only a small fraction of the neutrons generated will actually get into the core. If we assume a compressed core diameter of 6 cm, and a target-to-core distance of 30 cm (remember, it has to be safety outside the implosion system!), then only about 3% of the neutron flux will enter the core - an arrival rate of 15,000 neutrons/nanosecond using a 10 amp ion source. This many neutrons will significantly accelerate the chain reaction, cutting it by some 15 multiplication intervals.

The ENI does not have to be placed near the actual fission assembly. Since warhead dimensions are typically no more than 1-2 meters it can be placed virtually anywhere in the weapon, as long as there isn't a thick layer of moderating material (plastic, hydrocarbon fuel, etc.) between the ENI and the fission core.

The power supply required to drive a pulse tube has many similarities to the EBW pulse power supply. A pulse of a few hundred volts at a few hundred amps is needed to drive the ion source, and a 130-170 KeV pulse of several amps is required to extract the ions and accelerate the beam. This high voltage pulse controls actual neutron production and should thus have as fast an onset time as possible. This high voltage pulse can be supplied by discharging a capacitor of several KV through a pulse current transformer.

Pulse neutron tubes have been available commercially for decades (in non-miniaturized form) for use as a laboratory neutron source, or for non-destructive testing.

An additional type of ENI, not based on fusion reactions, has been successfully tested but apparently never deployed. This is the use of a compact betatron, a type of electron accelerator, to produce energetic photons (several MeV). These photons cause photon induced fission, and photon -> neutron reactions directly in the core.

4.1.8.3 Internal Tritium/Deuterium Initiators

Another approach to making an internal neutron initiator is to harness the high temperatures and densities achieved near the center of an implosion to trigger D+T fusion reactions. A few tenths of a gram of each isotope is placed in a small high pressure sphere at the center of the core in this scheme.

The number of actual fusions produced is small, but it may seem surprising that any could occur at all. The occurrence of fusion during a collision between two nuclei is a statistical process. The probability of it occurring on a given collision depends on the collision velocity. The velocity of the nuclei is in turn a statistical process which depends upon the temperature. The hydrogen plasma is in thermal equilibrium with a mean temperature of a few hundred thousand degrees K, but the Maxwellian energy distribution means that a very small number of ions is travelling at velocities very much higher than average. Given the very large number of ions present, a significant fusion rate results. Only a few fusions are actually necessary for reliable initiation after all.

The main attraction of this scheme is that the half-life of tritium (12.3 years) is much longer than Po-210, so the initiator can be stored ready-to-use for long periods of time. The system is also

physically simpler, and more compact than ENIs. It is not clear whether this type of initiator has actually been used in weapon designs.

4.1.9 Testing

Like any munition, the development of a fission weapon will require a variety of tests. These include component tests, and perhaps tests of the complete weapon. Tests of components like the firing system, detonators, etc. are similar to the requirements of non-nuclear munitions and need no comment. Even conservative gun assembly designs will normally require proof testing of the gun/propellant combination to verify the internal ballistics.

In addition to these routine types of tests, fission weapon development requires (or at least benefits greatly) from certain types of test that are unique to nuclear weapons. These include nuclear tests, by which I mean tests of the nuclear properties of materials and designs, not nuclear explosions (although an actual explosion of substantial yield is one possible type of nuclear test). Implosion designs, by which I mean any design using shock waves for core assembly, also call for hydrodynamic tests - tests of materials under the extreme conditions of shock compression. Combined nuclear and hydrodynamic tests, called hydronuclear tests, provide a more direct way of developing data for weapon design, evaluating design concepts, or evaluating actual designs. "Hydronuclear" is a somewhat vague term. Hydronuclear tests can mean shock compression experiments that create sub-critical conditions, or supercritical conditions with yields ranging from negligible all the way up to a substantial fraction of full weapon yield. Tests of negligible yield are often called "zero yield tests", although this is also not a precise term. Generally it is taken to mean a test in which the nuclear energy release is small compared to the conventional explosive energy used for assembly - a few kg of HE equivalent for example. However even in sub-critical tests the nuclear energy release is not actually "zero". It appears that the Comprehensive Test Ban Treaty (CTBT) now being negotiated in Geneva will use a "no-criticality" standard for defining legal experiments with high explosives and fissile material.

4.1.9.1 Nuclear Tests

A variety of nuclear tests are of interest for collecting design data. Since the performance of nuclear weapons is the combined effect of many individual nuclear properties, the most desirable measurements for weapon design purposes are "integral experiments" - experiments that directly measure overall weapon design parameters that combine many different effects.

Critical mass experiments determine the quantity of fissile material required criticality with a variety of fissile material compositions and densities, in various geometries, and with various reflector systems. These provide a basic reference for evaluating nuclear computer codes, estimating material requirements for weapons, and (extremely important) for doing safety evaluations. The closer the critical mass experiment resembles actual weapon configurations, the more useful it is. A considerable amount of critical mass data has been published openly which makes it possible to perform reasonably good "first cut" weapon design evaluations using scaling laws (like the efficiency equations). Any weapon development program will want to

perform criticality tests of systems closely resembling actual proposed designs, differing only in the amount of fissile material present.

Critical mass values can be predicted with good accuracy by extrapolation by taking neutron multiplication measurements in a succession of sub-critical tests using increasing quantities of fissile materials. Such tests can be conducted safety in the laboratory without special protective equipment since each successive test allows progressive refinements of critical mass estimates, and allows the calculation of safe masses for the next test. Tests intended to closely approach or reach criticality must be conducted under stringent safety conditions however. Even a very slight degree of criticality in an unmoderated system can produce a deadly radiation flux in seconds. Accidents during critical mass experiments killed two researchers at Los Alamos in 1945 and 1946 (Harry Daghlian and Louis Slotin) before manual experiments were banned there.

Basic critical mass tests are basically non-multiplying and do not measure alpha, the extremely important fast neutron multiplication parameter. Direct measurements of this require establishing systems with significant levels of supercriticality capable of creating rapid increases in neutron populations.

A variety of laboratory tests can be used for this. All of them depend on creating a supercritical state that persists for a very short period of time (milliseconds to microseconds) to prevent melt-down (or worse). Such experiments necessarily produce large neutron fluxes, and thus must be conducted under remote control.

One type of experiment creates a transient supercritical state by propelling a small fissile mass though a larger slightly sub-critical mass. The supercritical state exists while the small mass is inserted, and terminates when the mass exits the other side. Examples of this type of experiment are the "Dragon" experiments conducted at Los Alamos in early 1945, in which a fissile mass was dropped through a hole bored in a subcritical assembly (so-called because it was like "tickling the tail of a dragon"). Shorter assembly times (and thus higher multiplication rates) can be investigated by using a gun instead of gravity to accelerate the fissile projectile. This approach obviously extends naturally to evaluating a full-up gun weapon design with only the amount of fissile material in the bullet or target differing from the actual deployment weapon. This type of test was actually used by South Africa for evaluating its gun assembly weapon (using a test device named "Melba"). These experiments can explore assembly durations in the range of 0.1-10 milliseconds.

A second type of experiment achieves an even higher multiplication rate under controlled conditions by using the thermal expansion of the core to shut down the reaction. This is called a fast neutron pulse reactor. A solid core of fissile material is assembled that is slightly supercritical at room temperature, but is kept subcritical by the presence of a control rod, by removing a section of reflector, or by controlling the insertion of fissile material. When the rod is removed (or reflector is inserted), it becomes supercritical and rapidly heats up. The expansion of the material at sonic velocity, mediated by an acoustic wave, shuts down the reaction in a matter of microseconds. Assembly durations of 5-500 microseconds can be investigated.

Examples of this type of experiment are a series of fast pulsed reactors operated by the US during the late forties and early fifties: the bare uranium core Godiva, the bare plutonium core Jezebel, and the reflected uranium and plutonium assemblies Topsy and Popsy.

These fast metal assemblies can also be used to collect multiplication data at the border of criticality by adjusting their density in various ways. All of the mentioned US assemblies have been used to measure multiplication rates by studying the change in rates with density in the region between delayed and prompt criticality. These measurements can be extrapolated to estimate the maximum values of the materials. Although little data on alpha values for weapons-usable material have been published in general, results of these types of experiments are available.

4.1.9.2 Hydrodynamic Tests

Hydrodynamic tests can evaluate shock compression techniques and designs, and collect data on the properties of nuclear materials under shock compression conditions. The latter sort of test requires conducting shock experiments with actual nuclear materials of course.

This is not much of a problem from a safety point of view for uranium since comparatively non-toxic and nuclearly inert natural or depleted uranium is available. Hydrodynamic tests on complete implosion weapon designs can be conducted for uranium weapons simply by substituting natural uranium or DU for the actual U-235 or U-233.

This is not true for plutonium. There is no non-toxic, non-fissile form of plutonium. The radiotoxicity of plutonium make hydrodynamic tests much more hazardous to perform and care to avoid criticality is essential. It is interesting to note that a considerable amount of high-pressure shock equation of state data has been published for uranium, but very little or none has been for plutonium. Uranium can be used as a plutonium substitute to some extent, but the unique and bizarre physical state diagram of plutonium limits this to some extent. This is especially true in situations were very accurate EOS knowledge is required. The small safety margins involved in creating one-point safe sealed pit weapons, and in preparing for hydronuclear tests, places a premium on precise knowledge of plutonium behavior.

Measurements in weapon-type implosion systems are very difficult to make since they must be taken through the layer of expanding explosion gases. Flying plate systems are widely used for collecting equation of state data ranging up to fairly high shock pressures (several megabars). Advanced weapon programs typically use sophisticated instruments like light gas guns to generate very high pressure shock data.

Even with natural uranium or DU, full scale hydrodynamic test of weapon designs will require special test facilities including heavily reinforced test cells, with provision for instrumentation . The cells will unavoidably remain contaminated with detectable levels of uranium, showing the nature of tests that have been conducted there.

4.1.9.3 Hydronuclear Tests

Hydronuclear tests are the ultimate in integral experiments, since they combine the full range of hydrodynamic and nuclear effects. Although implosion weapons (i.e. Fat Man) have been successfully developed without any tests of this kind, a weapon development program is likely to regard such tests as highly desirable.

In hydronuclear tests of a candidate weapon design data on both the rate of increase of alpha during compression, and the maximum alpha value achieved can be collected. The first type of data is useful to determine the ideal moment of initiation for maximum efficiency, the second for determining how efficient the weapon will be.

The influence of time absorption and other effects dependent on neutron energy (fission cross sections, moderation, inelastic scattering, etc.) changes with effective multiplication rate. This encourages weapon developers to conduct tests at very high multiplication rates to collect good data for weapon performance prediction. Since weapon efficiency and yield are dependent primarily on the effective multiplication rate, this means tests with large releases of nuclear energy. Prohibiting tests with substantial nuclear energy yields (tens or hundreds of tons) may not prevent a nation from developing fission weapons, but it does at least restrict its ability to predict weapon yield.

A serious problem with hydronuclear tests is predicting what is going to happen in advance. On one hand, it is obvious that if one can predict *exactly* what will happen, then there is no need for the test at all. On the other hand, not being able to estimate the effects reasonably well in advance makes conducting the test extremely difficult, even perilous.

The reason for this should be clear from the efficiency equations. Since at low degrees of supercriticality efficiency and yield scale as (rho - 1)^3, fairly small variations in compression cause fairly large variations in yield. For example, if a two-fold compression factor is intended in create a supercritical density of rho = 1.02 (and a yield of say, 50 kg), then a 5% variation in compression could cause a result ranging from a complete failure to approach criticality, to a 45-fold overshoot (2.2 tonnes). Since designing suitable instrumentation requires having a fairly good knowledge about the range of conditions to be measured, the first would result in no data been collected. The second could destroy the test facility (and also result in no data being collected!). Actual US tests have been known to overshoot target yields of kilograms, producing yields in the tens and even hundreds of tons.

4.2 Fission Weapon Designs

The design elements outlined in the previous subsections can be applied in various ways to achieve a wide variety of weapon designs. The most appropriate design depends on the technology and material constraints within which the designer works, and what the design objective is (maximum efficiency, light weight, warhead volume and dimension restrictions, ruggedness, safety, etc.). The available fissile material is one of the most basic constraints in which there may be little flexibility. In this subsection discuss some particular design approaches that have been taken (or might have been, since a lot of this is speculation) for various purposes.

4.2.1 Low Technology Designs

These are designs that rely on a minimum of exotic technology or lengthy and costly research programs. Reasons for using such a design include: limited resources, low cost, low risk, and speed of development.

Situations where these issues might arise include:

1. A non-industrialized or partially-industrialized nation attempting weapon development;
2. A country undertaking a crash program due to sudden internal or external political changes; or
3. A country undertaking a weapon program with very limited objectives (such as merely demonstrating that they have nuclear capability).

There is the obstacle that production of weapon-grade fissile material is inherently a rather exotic and expensive process. I assume that suitable fissile material is at hand and thus only address actual weapon design issues.

There are a couple of well documented weapon designs that used nothing that would be considered particularly exotic or at all high-tech today: the first atomic bombs used in WWII - Little Boy, and Fat Man. Both of these devices are described in some detail in Section 7.1. Other weapons that have been fairly well documented are quite similar to these: the first Soviet atomic bomb (actually a copy of Fat Man), the first British atomic bomb (a Fat Man type weapon with modest improvements), and the South African gun-type atomic bomb.

4.2.1.1 Gun Designs

These are without a doubt the easiest weapons to manufacture if sufficient HEU is available. Mundane ordnance technology suffices for the assembly mechanism, standard industrial techniques are sufficient for fabricating a tungsten or tungsten carbide reflector, and a neutron initiator is neither difficult (if a modest reactor is at hand for polonium production) or even necessary.

The Little Boy design was completed (except for the fissile material) in a matter of months, and no need was felt to test it before combat use.

The South African nuclear weapon program is an example of Situation 3 - a limited objective program. South Africa intended to use their weapon as a political tool, not a military one. Due to the ease of development and inherent reliability a gun design was chosen.

4.2.1.2 Implosion Designs

Implosion weapons are more efficient, but are inherently more complex. The two principal problems are development of an effective implosion system and a neutron initiator. The basic techniques and designs for both are well known today, but actually fabricating working components (and being sure they work!) is a different matter entirely.

With the advantages of modern commercial technology for manufacturing, instrumentation, and computing; of a well developed scientific and engineering base in the open literature covering relevant technologies; and with successful approaches already known, the task of converting these approaches to proven hardware should take dramatically less effort than they did in 1944-45 (a few thousand person-years).

The decision of the USSR and UK to copy the Fat Man design for their first weapons indicates a mix of Situations 2 and 3 to varying degrees. Both programs were rushed, and attempted to simply demonstrate a capability as their first objective. Iraq's crash effort to produce a single weapon during the Gulf War probably relied on a simple implosion design for the same reasons.

4.2.2 High Efficiency Weapons

Most nuclear weapons programs reach a phase where using scarce and expensive fissile material most efficiently becomes a high priority (if fissile material eventually becomes abundant this may be replaced by other optimization criteria).

The principal means of achieving high efficiency in a pure fission design is to maximize the compression of the fissile material by concentrating a great deal of compressive energy in the core as uniformly as possible. This implies the use of a large mass of high explosives, and a levitated pit design, probably together with flying plate techniques.

The most likely approach is to use a pit within a larger hollow shell that acts as a driver. The high explosive layer surrounding the driver shell accelerates it inward, a sufficiently high HE/driver mass ratio being used to achieve velocities approaching the theoretical limit (8 km/sec). As the driver collapses and thickens, additional concentration of energy on the inner surface may occur. This will create the strongest possible shock when it impacts on the pit, the use of a buffer layer (beryllium perhaps) surrounding the outer shell of the pit may reduce the energy lost to entropy. The outer pit then implodes on the inner levitated portion of the pit.

This does indeed appear to be the design approach taken in the most efficient pure fission device ever tested by the US - the Hamlet device designed by Theodore Taylor, and detonated in Upshot-Knothole Harry (19 May 1953). This device used a 60 inch in diameter implosion system (based on the TX-13D bomb then in development) that weighed 7000 pounds (8000 pounds was the total device weight). This has been described as a test of a "hollow" design, but since the Type D levitated pit used in Hamlet had been previously tested in other weapons, and hollow levitated cores had been demonstrated years before, this may refer to the use of a hollow driver shell. The yield of Hamlet was predicted to be 37 kt, the actual yield was 32 kt.

Other techniques that are valuable:

- The use of an efficient tamper and reflector (preferably non-moderating);
- The use of a composite core (if two different fissile materials are available);
- Making the weapon a high yield design, which are inherently efficient (see below); and
- Fusion boosting, which takes us out of the realm of pure fission designs.

4.2.3 Low Yield Weapons

This represents one extreme of the weapon design spectrum - nuclear devices intend to make "small" explosions. Low yield in this context generally means yields much less than the 20 kts of a nominal fission weapon - say, 1-1000 tons. These are, of course, always very large compared to any other types of weapon of remotely similar size. They are small only in comparison to the potential capabilities of nuclear weapons.

The smallest nuclear weapons actually deployed have had yields around 10 tons (like the W54), and have been intended for short range tactical or nuclear demolition use (e.g. blowing up roads and bridges).

A low yield weapon can be made simply by taking an existing weapon and reducing its efficiency in some manner - like reducing the amount of explosive to create a weak implosion. But this likely result in a low-yield the weapon with unnecessarily high mass, volume, and cost.

A weapon designer will probably want to optimize a low yield weapon toward one of two design goals: minimizing its size or minimizing its cost (basically this means minimizing the fissile content of the device). Real weapons typically try to strike a balance between the two extremes.

4.2.3.1 Minimum Size

A low yield minimum mass or volume weapon would use an efficient fissile material (plutonium or U-233), a low mass implosion system (i.e. a relatively weak one), and a thin beryllium reflector (thickness no more than the core radius). Since volume increases with the cube of the radius, a thick layer of anything (explosive or reflector) surrounding the fissile core will add much more mass than that of the core itself.

Referring to the Reflector Savings Table 4.1.7.3.2.2-3 we can see that for beryllium thicknesses of a few centimeters, the radius of a plutonium core is reduced by 40-60% of the reflector thickness. Since the density difference between these materials is on the order of 10:1, substantial mass savings can be achieved. At some point though increasing the thickness of the reflector begins to add more mass than it saves, this marks the point of minimum total mass for the reflector/core system.

In general, minimum mass and minimum volume designs closely resemble each other. The use of a hollow core adds negligibly to the overall volume.

At the low end of this yield range (tens of tons) simply inducing the delta -> alpha phase transition in a metastable plutonium alloy may provide sufficient reactivity insertion. In this case a classical implosion system is not even necessary, a variety of mechanisms could be used to produce the weak 10-20 kilobar shock required to collapse the crystal structure.

Since the fissile core would be lightly reflected, and weakly compressed, a relatively large amount of fissile material is required: perhaps 10 kg for even a very low yield bomb. The efficiency is of course extremely poor, and the cost relatively high.

The absolute minimum possible mass for a bomb is determined by the smallest critical mass that will produce a significant yield. Since the critical mass for alpha-phase plutonium is 10.5 kg, and an additional 20-25% of mass is needed to make a significant explosion, this implies 13 kg or so. A thin beryllium reflector will reduce this, but the necessary high explosive and packaging will add mass, so the true absolute minimum probably lies in the range of 10-15 kg.

The W54 warhead used in the Davy Crockett had a minimum mass of about 23 kg, and had yields ranging from 10 tons up to 1 kt in various mods (probably achieved by varying the fissile content). The warhead was basically egg-shaped with the minor axis of 27.3 cm and a major axis of 40 cm. The W-54 probably represents a near minimum diameter for a spherical implosion device (the U.S. has conducted tests of a 25.4 cm implosion system however).

The test devices for this design fired in Hardtack Phase II (shots Hamilton and Humboldt on 15 October and 29 October 1958) weighed only 16 kg, impressively close to the minimum mass estimated above. These devices were 28 cm by 30 cm, Humboldt used PBX-9404 as the explosive.

4.2.3.2 Minimum Fissile Content

The contrasting approach to minimizing size is to make a small explosion in the most efficient way as possible. This means applying the same principles as high efficiency design, but simply reducing the amount of fissile material to reduce the yield. The mass of the implosion system, and the tamper/reflector in this case will result in greater overall mass and volume, even though the fissile material weight is reduced.

Using an advanced flying plate design it is possible to compress a 1 kg plutonium mass sufficiently to produce a yield in the 100 ton range. This design has an important implication on the type of fissile material that can be used. The high compression implies fast insertion times, while the low mass implies a low Pu-240 content. Taken together this means that a much higher Pu-240 content than normal weapon grade plutonium could be used in this type of design without affecting performance. In fact ordinary reactor grade plutonium would be as effective as weapon grade material for this use. Fusion boosting could produce yields exceeding 1 kt with this system.

4.2.4 High Yield Weapons

A nominal yield fission weapon uses one critical mass of material (at normal density) and has a yield around 20 kt. HEU has a larger critical mass than plutonium, but its efficiency is lower so the yield of a nominal weapon of either material is roughly the same.

High yield fission weapons use more than one critical mass of material. These weapons necessarily use hollow core designs, since this is the only way to render the core subcritical. High yield designs are inherently more efficient than nominal designs (assuming complete assembly occurs) since the large core radius reduces neutron leakage, and takes longer to disassemble. The first factor experiences diminishing returns as the core size grows and leakage becomes small, eventually becoming negligible for the core as a whole. For this reason

reflectors have little value in high yield designs, although by reducing the drop in neutron flux near the surface they help fission this outer layer more efficiently. The second factor (longer disassembly time) continues to enhance efficiency regardless of how large the core becomes, eventually though other factors begin to limit efficiency (see below). Tampers assist in retarding disassembly in high yield designs and probably significantly increase efficiency regardless of size. This is because they reduce the loss of the outer layers of material early in disassembly, allowing more of this material to fission.

A high yield core becomes critical comparatively early in the implosion process, perhaps before the imploding shell has even impacted on the levitated core. This means the period during which predetonation can occur is much longer. This considerably limits the usefulness of plutonium in a high yield bomb, since large masses also mean higher neutron emission rates. If the amount of explosive is limited, the large core implodes at a significantly slower rate as well.

A plutonium bomb similar to the Fat Man design, but containing four times as much fissile material (25 kg) would have a core diameter of 18 cm. To implode to the same final density (about 40) at the same velocity (2 km/sec) would take 18.7 microseconds, 4 times as long. The very low Pu-240 content of the plutonium produced during WWII (0.9%) would still give a reasonable chance of complete assembly but more economical grades (with higher Pu-240 content) would not. Such a design would have a yield in excess of 100 kt. The limiting efficiency of ~50% (see below) would give a yield of 210 kt. Higher implosion velocities are possible (permitting higher probabilities of optimum yield, or cheaper grades of plutonium), but this gives an indication of the practical limit for high yield plutonium fission bombs.

At a time when France had no access to enriched uranium, and had not yet developed fusion boosting technology, they developed plutonium bombs with yields of up to 120 kt (the MR31 missile warhead), probably the highest yield pure plutonium, pure fission device ever developed. The plutonium grades produced by the French had considerably lower burnups than US weapon grade plutonium (up 7% Pu-240), usually around 2% Pu-240, although "super-super-grade" plutonium (like the WWII US material) could have been produced especially for this weapon.

HEU can be used to make much larger weapons than plutonium due to its very low neutron emission rate. HEU pure fission weapons exceeding 1 megaton are possible. In very large fission bombs (hundreds of kilotons) the major disadvantage of HEU, its lower maximum alpha, disappears. This is because the race between the exponential growth in energy release and the disassembly of the core stops being the limiting factor in efficiency. Instead the problem of dilution of the fissile material by the fission products comes into play as the limiting factor. This limits efficiency to a maximum of about 50%.

An additional advantage in using HEU in large fission bombs is its cheapness relative to Pu-239 and U-233.

The largest pure fission bomb ever tested was the Mk 18F Super Oralloy Bomb (SOB) designed under the leadership of Dr. Theodore B. Taylor at Los Alamos. It demonstrated a yield of 500 kt

in the Ivy King test at Eniwetok (15 November 1952 local). Predicted yield was 400-600 kt. 85% of the yield came from U-235 fission, the rest presumably from fission of a U-238 tamper. This bomb used the large diameter (60 inch) 92 point implosion system developed for the Mk 13 high yield fission bomb, and the Mk 6 bomb casing and components. The Mk 18 weighed 8600 lb, about 90 were eventually deployed.

A reasonable assessment of the Mk 18 design is that it had a core containing 75 kg of HEU with a pre-implosion diameter of at least 24 cm, the levitated pit probably had a mass of 15 kg or so. It likely had a natural uranium tamper weighing about 150 kg. A density increase over the normal value of 2-2.5 is probable.

Safety is a serious problem with high yield fission bombs. Since several critical masses are present, simply collapsing the hollow space inside the core can render it highly supercritical. This does not require accurate implosion. Any accidental detonation of the explosive layer would squash a hollow core like a stomped tennis ball, and could lead to a very powerful explosion (in the tens of kilotons). Much milder accidents could also create serious criticality events. For example the possible breakage of a levitated pit support, allowing the levitated core to fall onto the hollow fissile shell. This does not change the overall density of the core, but it could create a local region of criticality where the levitated sphere rested on the layer of fissile material. Four approaches are available to reduce these problems:

1. Keep the bomb core partially disassembled, with the fissile material brought into its "implosion ready" configuration shortly before detonation.
2. Fill the hollow core with something that will prevent its collapse, then remove the material as part of the arming sequence.
3. Fill the cavity with a good fast neutron absorber (i.e. something containing boron-10) to provide an additional margin of criticality safety.
4. Insert a continuous neutron emitter of sufficient strength to guarantee early predetonation.

All four of these methods can be used together. The Mk 18 used the first three techniques, while the British high yield devices (such as Orange Herald) also used a removable neutron emitter.

Like several other pure fission bombs designed after the war, the Mk 18 used an automatic in-flight insertion mechanism to assemble the core. Just as the Gadget was assembled prior to the Trinity test by inserting part of the pit and the covering explosive lenses by hand, a motor was used to insert part of the fissile shell high explosive layers.

To prevent collapse of the core, and to enhance the very marginal degree of criticality safety, chains made of boral (boron-aluminum alloy) were inserted in the core. These chains had the problem that they could not be reinserted once removed. British high yield designs used ball bearings for this purpose, which were drained out the bottom during arming (these offer the possibility of refilling from the top). Heavy inert liquids have also been proposed.

The British used neutron emitters as a safety in the Orange Herald test device, but this was never adopted in the US.

4.2.5 Special Purpose Applications

4.2.5.1 Thermonuclear Primaries (Triggers)

Multi-stage thermonuclear weapons require as their first stage a fission bomb primary or trigger. The primary functions as part of a system to create the conditions for thermonuclear energy release. There are a number of possible primary design variations, partly due to the different design approaches that exist for thermonuclear weapons. Thermonuclear weapon and primary design is discussed in Section 4.4, but the essential feature of a suitable trigger is easily stated: energy must escape from the fissile core into the radiation case (surrounding the primary) very rapidly. This implies that the layers of material surrounding the fissile core must be transparent to the emitted thermal radiation. It is also desirable that the bulk of the radiation be emitted at a high temperature since the radiation implosion process is driven most efficiently by a high temperature photon gas. Rapid escape of the radiation also means that only a small fraction of the energy is deposited as kinetic energy in the trigger debris. This is very important because the impulse generated by debris collisions can potentially disrupt the implosion process.

In Subsection 4.1.2.7 (Post Disassembly Expansion) the progress of the expanding shock from the core is described until it reaches low-Z material outside the tamper, like a beryllium reflector or a high explosive containing no elements heavier than oxygen. If the mass of this low-Z material is not too great, then it will quickly become completely ionized and transparent. The high temperature bomb core will not reach thermal equilibrium with the reflector and explosive layers, and radiation will escape through them without further substantial heating.

The ionization of the outer layers of the primary, and the subsequent radiation cooling of the core can be significant while the fission energy release is still going on. In this case the cooling of the core surface delays expansion and contributes to enhancing primary efficiency.

Once the core (and tamper, if present) begin expanding it quickly forms a thin shell of high density, high-Z material which radiates away the trigger's thermal energy into the bomb casing. Nearly all of this energy exists as a photon gas with a uniform high temperature in the low density region inside the expanding shell. A temperature gradient quickly becomes established across the shell thickness as this energy quickly flows from the interior of the fireball into the radiation case.

If the mass of material outside of the core and tamper is not small however, and worse still, also contains significant amounts of high Z material then this process of energy transport out of the core is not efficient. Instead the core reaches thermal equilibrium with the reflector and explosive, diluting the thermal energy with the large mass of material. The thermal energy diffuses out of the opaque mass relatively slowly, and a large percentage of the energy is converted into kinetic energy in the primary debris.

The original Fat Man design is good example of a poor trigger. A thick layer of high explosive surrounded the tamper and core, and this explosive contained large amounts of barium, a relatively high-Z material, due to the use of baratol as the slow explosive component of the lens. The explosive energy of the core, amounting to some 20 kt was diluted by about 2500 kg of HE in a volume 140 cm across before it could escape into the radiation channel.

The triggers of the earlier thermonuclear devices like the Sausage (Ivy Mike test) and Shrimp (Castle Bravo test) were similar to the Fat Man system, but had thinner, less massive explosive lens systems (100 cm across, 1000 kg of HE) due in part to the use of a larger number of explosives lenses (92 vs 32). But the most important difference was the use of boracitol instead of baratol, eliminating any atom with a higher Z than oxygen. The thermal radiation emitted by the core was thus able to completely ionize the explosive layer, rendering it transparent, and allowing the rest of the energy to escape the core unimpeded.

Modern boosted fission triggers take this evolution towards higher yield to weight, smaller volume, and greater ease of radiation escape to an extreme. Comparable explosive yields are produced by a core consisting of 3.5-4.5 kg of plutonium, 5-6 kg of beryllium reflector, and some 20 kilograms of high explosive containing essentially no high-Z material. Explosives lenses incorporating boracitol or inert filled plastic foams may be used or, more likely, the classical explosive lenses may have been replaced other advanced wave shaping techniques.

Light weight primaries of this type invariably use fusion boosting (see Subsection 4.3.1) to compensate for the limited degree of reactivity insertion that can be achieved with such small amounts of explosive and fissile material.

In these triggers, thermal radiation escaping from the core completely ionizes the low-Z beryllium and the explosive layers, even before the core disassembles (that is - while the fission reaction is still underway). The approximately 100-fold improvement in yield to mass ratio over Fat Man leads to a similar increase in achievable radiation density inside the bomb casing (and a greater than three-fold increase in temperature).

Within these general design guidelines, significantly different types of primaries can still be developed (discussed further in Subsection 4.4 Elements of Thermonuclear Weapon Design).

4.2.5.2 Earth Penetrating Warheads

The destruction of hardened underground structures (like command bunkers, missile silos, sub pens etc.) is much more efficient if the explosion occurs underground. Surface bursts and air bursts do not transmit energy efficiently to the ground, giving a moderate sized explosion a relatively small radius of effectiveness. An underground explosion, even a relatively shallow one, converts nearly all of its energy into a ground shock wave. If the warhead can burrow down to the same depth as the target, its effectiveness if enhanced even more because it is closer to its target. A shallow penetrating warhead produces very high levels of local fallout contamination, although a deep penetration warhead can potentially reduce the amount of radiation found at the surface.

Designing a weapon that can penetrate deeply into the ground (which may also be paved with thick reinforced concrete) is a significant problem. The basic requirement for a ground penetrator weapon is to have a nuclear device inside a long, narrow, strong casing that is massive and strong enough to punch through concrete, rocks, and soil. The nuclear device must also be rugged enough not to be damaged by the shock of impact. Different device designs could be packaged in casing to meet this requirement, but spherical implosion systems would need to have a small radius to fit inside such a package. Thermonuclear systems have been hardened to withstand accelerations at least up to 3000 Gs.

This is an application to which gun assembly weapons are uniquely suited. These weapons are intended for destroying (like ICBM silos, and control bunkers etc.)

The requirement for a long thin, heavy, very strong bomb nicely matches the physical features of a gun tube. During Desert Storm conventional HE "bunker busters" were made by pouring TNT into actual artillery barrels. Such a bomb can penetrate as much as 100m into the ground, and can punch through several meters of reinforced concrete in addition. As a result earth penetrating bombs have been the major application of gun-type weapons since the 1940s. The US deployed gun-type earth penetrating bombs such as the Mk-8 "Elsie" and the Mk-11.

One possible problem with an earth penetrating gun weapon is that essentially all of the fissile material remains unchanged after the explosion. The material from one gun-type bomb is sufficient to manufacture 3-4 implosion bombs. Since earth penetrating bombs are inevitably targeted on enemy territory, this means a potential future adversary now has access to several bomb's worth of fissile material. Even though this is distributed through a few thousand tons of radioactive rock, mining and extracting would probably be relatively easy compared to setting up production facilities for fissile material. Soil normally contains 1-3 ppm of uranium, so the weapon grade material (50 kg, say) would be diluted by only several kg of natural uranium from this source. This problem could be reduced by including at least an equal mass of U-238 denaturing material at the far end of the bomb from the target, shielded by a neutron absorber.

The current US ground penetrating warhead is the recently developed B61-11 bomb. It was designed by repackaging a B61-7 thermonuclear warhead (which was inherently shock resistant) in a heavy high strength steel bomb body with a special nose. The depth of penetration is shallow (~6 meters).

4.2.6 Weapon Design and Clandestine Proliferation

A subject of considerable concern throughout the world is the problem of secret acquisition of nuclear weapons. Numerous treaties and agreements have been made between the nations of the world to restrict the spread and deployment of nuclear weapons, and to formalize the resolve by many nations not to acquire them. Even if all nations were to join these protocols, the concern would persist that one or more nations might break their agreement in order to gain an advantage over their neighbors. A number of regulatory regimes have been set up to circumvent this, but none of them are fool-proof.

The question arises then, what sorts of weapons could be acquired secretly or illicitly? International regulation and approbation towards proliferators impose significant restrictions on the types of activities in which a non-nuclear nation can engage in an effort to acquire weapons. The nuclear weapon states had the advantage of being able to pursue whatever weapons designs or technologies they wished, and to build the best weapons they knew how to build, in whatever numbers they felt they could afford. For nations seeking to develop weapons in secret, the technologies and industrial capabilities employed by the nuclear weapon states will be lacking in whole or in part. Design and development decisions must be made to circumvent these disadvantages.

Of course fears of clandestine proliferation is not limited to what a nation-state might do. There is reasonable ground to believe that any nation or national leader can be deterred from using weapons of mass-destruction, since the nation itself is also a hostage to others. This assurance does not exist if the party possessing these weapons is not a nation, and in fact has no identifiable address. Sub-national groups: rebels, guerrillas, criminals, and especially terrorists, are also cause for grave concern.

4.2.6.1 Clandestine Weapons Development and Testing

For a clandestine program the principal problem is acquiring weapon-usable fissile material. Without this, no program is possible. The options are to manufacture it, or to acquire it ready-made.

Manufacturing weapon-usable material is by any means or measure a very expensive and difficult proposition. Even the lowest cost options for limited amounts of material (construction of a 20 MW breeder reactor, and plutonium processing plant) run to US$100 million or more (possibly much more), and require establishing an entire industry.

A number of different paths have been chosen by nations attempting to secretly manufacture weapons-usable fissile material: plutonium reactors (Israel, India, North Korea, and initially Iraq), gaseous diffusion plants (Argentina), gas centrifuges (Pakistan, Brazil, Iraq, India), aerodynamic separation (South Africa), laser isotope separation (Israel) and even calutrons (Iraq). The choice is largely determined by the resources and technical capabilities of the nation, and the peculiar advantages that the nation may be able to secure.

By "peculiar advantages", I mean such things as technical assistance from outside nations (from France to Israel; collaboration between Israel and South Africa; from China to Pakistan), the ability to disguise it as or within a civilian program (Brazil, South Africa, Iraq, India), or the ability to steal or illegally divert technology and materials (Pakistan, Iraq, Israel, India, Brazil). Most programs depend on a combination of these things.

Both highly enriched uranium and plutonium have been acquired in secret programs. But so far apparently no nation has acquired both. This simply indicates that the cost and effort is so great that a nation is likely to successfully pursue only one option.

The chief problem with acquiring weapons usable material ready-made is finding a supplier. This is perhaps not as big a problem for a nation engaging in clandestine proliferation as it might appear. There has been a long-established international trade in weapons usable material. Highly enriched uranium is used in certain types of research reactors, in naval propulsion reactors, and in certain prototype power plants. Small amounts of weapons grade plutonium have been exchanged for research purposes as well. The civilian production and trade in reactor grade plutonium is already large and growing.

No clandestine weapons program is known to have actually made use of acquired weapon-usable fissile material, but only just barely. Iraq actually initiated a crash program during the 1990-91 Gulf War to seize the HEU from an internationally safe guarded research reactor in Iraq, and build a single nuclear weapon with it. This effort never got so far as to actually begin making the weapon, although Iraq did abscond with the fuel rods and tampered with them.

Of the three classes of weapon-usable material available for illicit acquisition - HEU, weapon grade plutonium, and non-weapon grade plutonium - the scarcest and mostly heavily protected is weapon grade plutonium. While the possibility of theft occurring in the territories of the former Soviet Union remains a concern, this material is relatively unlikely to be obtained in sufficient quantity by anyone, with or without the permission of its owner. HEU has been available for civilian uses for many years, and has not always been well guarded. Recent efforts have been undertaken by the US to eliminate the civilian use of HEU, but it remains a serious cause for concern. The rapidly expanding civilian production and use of non-weapon grade plutonium is the major concern for the present and the future. Many nations have or soon will have substantial stockpiles of this material, and will be able to divert significant quantities secretly. Secret diversion could then be followed by an abrupt, open, large scale breakout of the regulatory regime once a weapon design has been perfected. Numerous stockpiles and a large volume trade in plutonium greatly increase the risk of theft. Special attention must thus be paid to the possibilities of using reactor grade plutonium in clandestine weapons.

Once weapons-usable material has been acquired, actually designing and manufacturing weapons is the next issue. Compared to the problem of manufacturing fissile material, this is comparatively easy however. The fundamental technologies to actually build a weapon is possessed by any nation with a significant arms industry (that is, virtually any country with a significant military). The technologies used to actually build the weapons employed by the US in WWII are crude by today's standards, and are widely available.

Some desirable technologies used in advanced weapons are restricted in their availability. A famous example of this are the krytron high-speed switches that were illicitly sought by Iraq, Israel, and Pakistan. Miniature pulse neutron tubes, high purity beryllium and beryllium fabricating equipment, and advanced wave shaping technologies are other examples. But none of these are actually necessary to manufacture weapons.

These technologies are especially valuable for minimizing weapon weight, which is a key consideration if the weapon is intended to arm a ballistic missile. Ballistic missiles typically have a payload weight limit in the range of 500-1000 kg. However, since South Africa actually

manufactured a gun-type weapon with a weight of only 1000 kg, it is likely that this constraint is not too severe.

At one time computers, useful for numerical simulation in weapon design, were considered to be a restricted technology that limited the ability of other nations to develop weapons. However, this capability is most important for thermonuclear weapon design, not fission weapons. The computational effort required for the neutronic and hydrodynamic computations used in fission weapons is actually quite modest, easily within the capability of any commercial PC available today. Even with thermonuclear weapon design, computational requirements are not that extreme. The initial design effort on most weapons in the US arsenal (perhaps all of them) were completed before the first Cray 1 went on line in 1976. A high end workstation is comparable or superior to the best computers available when most current US warheads were developed. Even the lowest performance office computers now on the market are orders of magnitude faster than the computers that were used to design the first hydrogen bombs.

Of course raw computational power is not sufficient. Sophisticated codes, and extensive physical data are required to make use of them. Also the connection between weapon simulation and testing should not be forgotten. Sophisticated simulation capabilities permit nuclear weapon states to reduce or even eliminate the need for weapon tests to develop or prove a design. A country without an experience base in weapon design is at a significant disadvantage here. The lack of proven codes will substantially constrain the usefulness of computer technology.

A clandestine weapon developer will presumably use nuclear, hydrodynamic, and hydronuclear tests to the greatest practical extent as a substitute for full scale weapon tests. Hydronuclear testing is likely to be fairly expensive though, since great caution will be needed to avoid large yield overshoots that disclose the program prematurely. This will require numerous tests with scarce and costly fissile material to creep up to the desired test yield range. The fissile material can be reclaimed after each test of course (as long as it is not accidentally dispersed in an overshoot), but this takes time and either requires a large inventory of fissile material, or a very slow test program.

Since every nation to develop nuclear weapons appears to have succeeded on their very first weapons test (with the possible exception of India, the information here is conflicting), and other nations have deemed it unnecessary to even test their arsenal in advance (the US with Little Boy; South Africa and Pakistan), there is legitimate grounds for doubting how essential full yield testing really is. The US did not experience its first test failure until 1951 with its 18th test. It is clear that weapons can be built without full yield testing (or even hydronuclear testing), but considerable information can be obtained with sub-yield testing - especially for nations without prior test experience. In the absence of testing a nation will be forced to make use of more conservative, and less highly optimized designs, and will have a higher level of uncertainty about actual weapon performance. Certain design options (perhaps fusion boosting, as an example) may also be infeasible.

While simply copying the Manhattan Project weapon designs will provide a nation with a workable weapon, it is very likely that any nation developing weapons today will seek to improve upon them (unless the nation feels itself under pressure to produce a weapon as quickly as possible). South Africa produced a very high reliability (and presumably highly safe) gun-type weapon weighing 1000 kg, compared to 4500 kg for the Little Boy design. Nations developing implosion weapons today will probably attempt to make much lighter lens systems than the Fat Man design, and employ levitated pits, even on their first weapon design.

4.2.6.2 Terrorist Bombs

The prospect of terrorist acquisition of weapons has haunted the world since at least the late sixties, when international terrorism gained prominence. A variety of opinions have been expressed on the plausibility of these threats. Claims have been made that a terrorist weapon could:

- Produce a yield of many kilotons;
- Be made with reactor grade plutonium, perhaps even with unrefined plutonium reactor fuel;
- Be made by a single individual in a matter of weeks, with commonly available materials (given that fissile material is also available);
- Be made small enough to transport easily by car. These claims have even been conflated, so that it is asserted that most or all of them are true simultaneously.

These claims are all conditionally true: they may be valid, but only under a restrictive set of assumptions. And they also conflict strongly. Some are completely incompatible; others cannot be categorically eliminated as impossible in combination, but in any event it seems that no more two of them could be possible under any scenario.

What technologies are plausible for terrorist use? And what types of weapons are reasonable threats?

The most fundamental constraint on a terrorist group is the type of fissile material that is available, and in what quantity. The key problem is to obtain any fissile material at all, a terrorist is not likely to have any choice in the matter as too what kind it is. The breakup of the Soviet Union has brought about a worrisome trade in fissile materials. A significant amount of weapons-usable material has turned up, although intelligence agencies running sting operations are the only actual "market" so far identified. The quantities of materials that have surfaced so far have not been nearly enough for a weapon (even if it were all pooled), but the quantities have been large enough to cause considerable concern. At the moment the most frightening problem is Russian naval fuel. This contains highly enriched uranium (even higher than standard weapons grade!), and has been poorly secured at some locations. It can be hoped that in the near future the former Soviet fissile stocks will be secured sufficiently to eliminate this as an immediate concern.

In the long run, the availability of plutonium through commercial reprocessing for use in mixed oxide fuel (MOX) for commercial power reactors represents the major risk. Over one hundred tons of plutonium have already been commercially separated (an amount that will soon exceed the world's total weapons-grade plutonium production). This material will be in the hands of many nations, who will likely not all be equally vigilant in protecting their fuel stocks.

Clearly the most serious scenario is if weapons-grade HEU can be obtained by a terrorist group. Due to the very low neutron emission rate, very low technology can produce a substantial probability of full insertion and high yield detonation.

A weapon constructed from 40 kg of 93.5% HEU, with a 10 cm tungsten carbide reflector would produce a full yield of >10 kt. The required assembly time for a 50% chance of complete assembly is some 48 milliseconds, equal to a velocity of only 9 m/sec. This can be achieved by simply dropping the bullet 4.4 meters! Crude gun-type arrangements, along the lines of the IRA's makeshift mortars could easily achieve velocities of 100 m/sec or more.

A gun-type weapon is not a major concern if plutonium is used. Such a device might actually produce explosive yields in the range of a few tons, but would not be significantly more destructive than conventional truck bombs. On the other hand explosive compression, required for higher yields, is much more difficult to arrange. At the very least it requires a substantial quantity of good quality high explosive - at least a few hundred kilograms (unless the design and construction is rather sophisticated).

Now and in the future, reactor grade plutonium appears to be the material most likely to be available to a terrorist group. Given the spontaneous fission rate, and the limited technology for rapid assembly, predetonation is a foregone conclusion. In this scenario the yield of the system is not determined by the actual compression capability of the implosion system. Instead it is the rate of insertion that controls efficiency and yield. Any bomb design must emphasize making the insertion rate at the moment of criticality as fast as possible. In any case, rho (the density at the moment of disassembly relative to critical density) is going to be fairly small. Still, if insertion rates approaching those of the Fat Man design can be achieved then yields in the hundred of tons are possible.

Despite hints to the contrary (for example Ted Taylor's comments in _The Curve of Binding Energy_ among others), it is not plausible that true spherical implosion systems can be developed by a terrorist group. The difficulties in designing and making a working lens system appears to be simply insurmountable. Unfortunately, a spherical implosion system does not seem to be required for reasonably fast insertion at low levels of compression.

Consider an implosion of a system that may be in one dimension (linear implosion), two dimensions (cylindrical implosion), or three dimensions (spherical implosion). If delta represents the change in system dimension (i.e. size - radius or length) along the axis or axes of compression in n dimensions (n equals 1, 2, or 3), then the compression C achieved by the implosion is:

$$C = (r_0/(r_0 - delta))^n$$

At very low degrees of compression, this is roughly equivalent to:

$$C = n*(delta/r_0) + 1$$

That is, the excess density C - 1 is roughly proportional to the dimensional reduction ratio and the number of axes of compression. Thus for a given compression velocity, the actual rate of density increase for 3-D compression is three times faster than 1-D compression, but only 50% faster than 2-D compression. These differences are significant, but not dramatic.

Developing linear and cylindrical implosion systems fast enough to produce a highly destructive terrorist bomb appears to be feasible. The flying plate line-charge approach is sufficiently simple, and testable, that a low resource group could develop a workable system. Even plane or cylindrical explosive lenses are not out of the question, although they are probably more difficult.

Illicitly obtained plutonium would most probably be in the form of plutonium oxide, possibly as mixed oxide fuel. If the material were purified oxide powder, then it could be used directly in a bomb design. Fuel material, fabricated or not, would require chemical separation. A group sophisticated enough to attempt chemical processing would probably go on to reduce the plutonium to metal which is much more desirable for bomb construction.

Since the density of plutonium oxide is much since lower than plutonium metal, considerably more plutonium in this form would be needed. How much would depend on how highly compacted the plutonium oxide was at the moment of criticality. Although the crystal density of PuO_2 is 11.4, the bulk density of unconsolidated oxide powder is only 3-4 (possibly even lower). To raise it as high as 5-6 would require compacting under substantial pressure.

The pressures generated by shock waves are much less efficient at compacting porous materials, compared to static pressures. This is due to the inherent strong entropic heating associated with large volume changes during shock compression. However the pressures in a strong high explosive shock (or generated by an explosive driven high velocity plate collision) are so high that densities approaching the theoretical crystal density are probably achievable. If it is assumed that a bomb builder could compress the powder to a density of 5 with moderate pressure, and that a density of 10 is achieved during implosion, then something like 50 kg of plutonium in the form of oxide would be required for a bomb without a reflector. Assuming a fairly good, readily available reflector (a few inches of iron or graphite), this could be reduced to 25-30 kg. Taking into account the explosive required, such a bomb (with a reflector) would be large - weighing on the order of a tonne.

Using plutonium metal would greatly reduce fissile material requirements, and lead to a much smaller bomb. A design might use the cylindrical collapse of a hollow ring of plutonium metal (as

the delta or alpha phase), or cylindrical compression of a solid delta-phase aluminum-plutonium alloy disk. No more than about 10 kg of plutonium would be required in such a design, if a reasonably good reflector were used. Such a weapon might weigh as little as 200 kg.

Given that the system will disassemble well before compression is complete, an accurate symmetrical implosion is not really a necessity. Simply imploding the fissile material at a high rate even if imperfectly (that is, without a true plane or cylindrical shock wave), could produce the necessary rapid compression. For this to work, the fissile material would have to be fairly close to critical at the beginning of the implosion since an imperfect implosion would create unacceptable distortions if the compression factor were very large. As noted earlier in the discussion on nuclear testing, manufacturing a device that is close to critical is extremely hazardous and itself requires substantial sophistication.

4.3 Fission-Fusion Hybrid Weapons

The first designs proposed for fusion bombs in the U.S. assumed that the heat from the fission trigger would ignite a self-sustaining fusion reaction in a mass of liquid deuterium adjacent to it. In the late 40s and early 50s improved calculations showed that this was impossible. The only fusion reaction achievable by simply heating the fuel with a fission bomb is the D- T reaction:

D + T -> He-4 (3.5 MeV) + n (14.1 MeV)

The naive approach to using this reaction - making a large explosion by igniting a large mass of D-T fuel mixture with a fission trigger - is prohibitively costly. Plutonium is a factor of 10 times cheaper per unit of energy released compared to D-T fuel, and HEU is 3-5 times cheaper still. Furthermore due to radioactive decay the tritium continuously disappears at a rate of 5.5% annually and must be replaced.

A number of weapon designs have been developed that use the D-T reaction in a variety of ways however. All of them depend on the highly energetic neutrons produced by the D-T reaction. Some of these designs use the neutrons to achieve significant fission yield enhancement, thus reducing the expenditure of fissile material for a given yield. Others exploit the neutrons directly as a weapon.

The fusion boosting and Alarm Clock/Layer Cake designs were pioneered by the US and USSR in the early 1950s. Neutron bombs were apparently not developed by either nation until the late 1960s or early 1970s.

4.3.1 Fusion Boosted Fission Weapons

Fusion boosting is a technique for increasing the efficiency of a small light weight fission bomb by introducing a modest amount of deuterium- tritium mixture (typically containing 2-3 g of tritium) inside the fission core. As the fission chain reaction proceeds and the core temperature

rises at some point the fusion reaction begins to occur at a significant rate. This reaction injects fusion neutrons into the core, causing the neutron population to rise faster than it would from fission alone (that is, the effective value of alpha increases).

The fusion neutrons are extremely energetic, seven times more energetic than an average fission neutron, which causes them to boost the overall alpha far out of proportion to their numbers. Is this due to several reasons:

1. Their high velocity creates the opposite of time absorption - time magnification.

2. When these energetic neutrons strike a fissile nucleus a much larger number of secondary neutrons are released (e.g. 4.6 vs 2.9 for Pu-239).

3. The fission cross section is larger in both absolute terms, and in proportion to scattering and capture cross sections.

Taking these factors into account, the maximum alpha value for plutonium (density 19.8) is some 8 times higher than for an average fission neutron (2.5x10^9 vs 3x10^8).

A sense of the potential contribution of fusion boosting can be gained by observing at 1.5 g of tritium (half an atom mole) will produce sufficient neutrons to fission 120 g of plutonium directly, and 660 g when the secondary neutrons are taken into account. This would release 11.6 kt of energy, and would by itself result in a 14.7% overall efficiency for a bomb containing 4.5 kg of plutonium (a typical small fission trigger). The fusion energy release is just 0.20 kt, less than 2% of the overall yield. Larger total yields and higher efficiency is possible of course, since this neglects the fission-only chain reaction required to ignite the fusion reaction in the first place and that fission multiplication would continue significantly beyond the fissions caused by the fusion induced secondaries.

The fusion reaction rate is proportional to the square of the density at a given temperature, so it is important for the fusion fuel density to be as high as possible. The higher the density achieved, the lower the temperature required to initiate boosting. Lower boosting initiation temperatures mean that less pre-boost fission is required, allowing lower alpha cores to be used.

High fusion fuel densities can be achieved by using fuel with a high initial density (highly compressed gas, liquid hydrogen, or lithium hydride), by efficient compression during implosion, or most likely by both.

Although liquid D-T was used in the first US boosting test (Greenhouse Item), this is not a practical approach due to the difficulty in achieving and maintaining cryogenic temperatures (especially considering that 3 g of tritium constitutes a heat source of approximately 1 watt).

US nuclear weapons are known to incorporate tritium as a high pressure gas, that is kept in a reservoir external to the core (probably a deuterium - tritium mixture). The gas is vented into the weapon core shortly before detonation as part of the arming sequence. Initial densities with a

room- temperature gas (even a very high pressure one) are substantially lower than liquid density. The external gas reservoir has the important advantage though that it allows the use of "sealed pit", a sealed plutonium core that does not need servicing. The tritium reservoir can be easily removed for repurification and replenishment (removing the He-3 decay product, and adding tritium to make up for the decay loss) without disturbing the weapon core.

A possible alternative the use of a high pressure gas reservoir is to store the gas in the form of a metal hydride powder, uranium hydride (UH3) for example. The hydrogen can be rapidly and efficiently released by heating the hydride to a high temperature - with a pyrotechnic or electrical heat source perhaps.

A problem with using hydrogen gas is that it reacts very rapidly with both uranium and plutonium to form solid hydrides (especially plutonium, the Pu-H reaction rate is hundreds of times higher than that of any other metal). The formation of hydrides is very undesirable for the boosting process since it dilutes the gas with high-Z material. This can be prevented by lining boost gas cavity with an impermeable material. Thin copper shells have been used for this purpose. Alternatively the injection of fusion fuel could simply be conducted immediately before detonation, reducing contact between the core and the hydrogen isotope mixture to no more than a few seconds.

Lithium hydrides achieve an atomic density of hydrogen that is about 50% higher than in the liquid state, and since the hydride is a (relatively) stable inert solid it is also easy to handle. A key disadvantage is that the hydride must be permanently incorporated into the core requiring complete core removal and disassembly to replenish and purify the tritium.

The ideal location for the boosting gas would seem to be in a cavity in the very center of the fissile mass, since this would maximize the probability of neutron capture, and the core temperature is also highest there. In a levitated core design, this would make the levitated core into a hollow sphere. This is not desirable from the viewpoint of efficient fissile material compression however since a rarefaction wave would be generated as soon as the shock reached the cavity wall.

An alternative is to place the boosting gas between the outer shell and the levitated pit. Here the collapsing thin shell would create multiple reflected shocks that would efficiently compress the gas to a thin very high density layer. There is evidence that US boosted primaries actually contain the boosting gas within the external shell rather than an inner levitated shell. The W-47 primary used a neutron absorbing safing wire that was withdrawn from the core during weapon arming, but still kept its end flush with the shell to form a gas-tight seal.

The conditions created by compressing the gas between the collapsing shell and levitated core are reminiscent of a recently reported shock compression experiment conducted at Lawrence Livermore in which liquid hydrogen was compressed the metallic state by the impact of a 7 km/sec gas gun driven plate. This experiment generated pressures of 1.4 megabars, and hydrogen densities nine times higher than liquid. The velocity of an imploding shell is more like 3 km/sec and the boost gas is at a lower initial density, still, the pressures that can be expected

are at least as high, so a similar hydrogen density (around 0.75 atom-moles/cm^3) may be achievable.

It is also possible to dispense with a levitated pit entirely and simply collapse a hollow sphere filled with boosting gas. Since the fissile shell would return to normal density early in the collapse, there does not seem to be any advantage in doing this.

Fusion boosting can also be used in gun-type weapons. The South Africans considered adding it to their fission bombs, which would have increased yield five-fold (from 20 kt to 100 kt). Since implosion does not occur in gun devices, it cannot contribute to fusion fuel compression. Instead some sort of piston arrangement might be used in which the kinetic energy of the bullet is harnessed by striking a static capsule.

The fusion fuel becomes completely ionized early in the fission process. Subsequent heating of the hydrogen ions then occurs as a two step process - thermal photons emitted by the core transfer energy to electrons in the boost plasma, which then transfer energy to the ions by repeated collisions. As long as this heating process dominates, the fusion fuel remains in thermal equilibrium with the core. As the temperature rises, the fusion fuel becomes increasingly transparent to the thermal radiation. The coupling is efficient up to around 10^7 K, after which the fuel intercepts a dwindling fraction of the photon flux (which is should still keep it in temperature equilibrium given the greatly increasing flux intensity).

The fusion process releases 80% of its energy as neutron kinetic energy, which immediately escapes from the fuel. The remaining 20% is deposited as kinetic energy carried by a helium-4 ion. This energy remains in the gas, and can potentially cause significant heating of the fuel. The question arises then whether the fusion fuel continues to remain in equilibrium with the core once thermonuclear burn becomes significant, or whether self- heating can boost the fuel to higher temperatures. This process could, in principal, cause the fusion fuel temperature to "run away" from the core temperature leading to much faster fuel burnup.

I have not resolved this question satisfactorily at present, but it may be that the fusion fuel will remain in equilibrium, rather than undergo a runaway burn. Most of the helium ion energy is actually transferred to the electrons in the plasma (80-90%), which then redistribute it to the deuterium and tritium ions, and to bremsstrahlung photons. The energy must be transferred to the ions before it is available for accelerating the fusion reaction, a process which must compete with photon emission. If the photon-electron coupling is sufficiently weak then the boost gas can still runaway from the core temperature, otherwise it will remain in thermal equilibrium.

Boosting effectively begins when the ions are hot enough to produce neutrons at a rate that is significant compared to the neutron production rate through fission alone. This causes the effective value of alpha in the core to increase leading to faster energy production and neutron multiplication. In the temperature range where boosting occurs, the D-T fusion rate increases very rapidly with temperature (modelled as an exponential or high order polynomial function), so the boosting effect quickly becomes stronger as the core temperature climbs.

At any particular moment the contribution to alpha enhancement from boosting is determined by the ratio between the rate of neutron increase due to fission spectrum neutron secondaries, and the rate of increase due to fusion neutron secondaries. The fission spectrum contribution is determined in turn by the unboosted fission spectrum value of alpha, and the fission spectrum neutron population in the core. The fusion contribution is determined by the fusion reaction rate, and the fusion neutron alpha value. To optimize yield this enhancement should be at a maximum just as disassembly begins.

The fusion reaction rate typically becomes significant at 20-30 million degrees K. This temperature is reached at very low efficiencies, when less than 1% of the fissile material has fissioned (corresponding to a yield in the range of hundreds of tons). Since implosion weapons can be designed that will achieve yields in this range even if neutrons are present a the moment of criticality, fusion boosting allows the manufacture of efficient weapons that are immune to predetonation. Elimination of this hazard is a very important advantage in using boosting. It appears that every weapon now in the U.S. arsenal is a boosted design.

4.3.2 Neutron Bombs ("Enhanced Radiation Weapons")

The design objective of the tactical neutron bombs developed in the 1960s and 70s was to create a low-yield, compact weapon that produced a lethal burst of neutrons. These neutrons can penetrate steel armor with relative ease, enabling the weapons to be effective against tanks and other armored vehicles which are otherwise highly resistant to the effects of nuclear weapons. A flux of several thousand rems were desired so that incapacitation of armored crews would be relatively rapid, with in several hours to a couple of days at most. In this exposure range death is inevitable. To minimize the effects of collateral damage, the effect of thermal radiation and blast outside the neutron kill radius, it was also very desirable to minimize the energy released in forms other than the neutron flux.

The means for generating this intense neutron flux is to ignite a quantity of deuterium-tritium fuel with a low yield fission explosion. It is essential however to avoid the absorption of those neutrons within the bomb, and especially to *prevent* the fusion-boosting effect on the trigger. The weapon must also fit inside an 8" diameter artillery shell.

An example of such a weapon is the US Mk 79-0 warhead for the XM-753 8" AFAP (artillery fired atomic projectile). This shell was 44 inches long and weighed 214 lb. The W-79-0 component was only about 37 cm long. The maximum yield of the W-70-0 was 1 kt, of which 0.75 kt was due to fusion, and 0.25 kt to fission.

It has been suggested by some that a neutron bomb is simply a variation of a boosted fission bomb, e.g. the fusion fuel is in the center of the fissile mass. Elementary analysis shows that this idea is impossible. The 3:1 fusion:fission yield ratio of the W-79-0 indicates that there must be 31 fusion reactions releasing 540 MeV (and 31 fusion neutrons) for each fission (which release 180 MeV). This means more than 97% of the fusion neutrons must escape the core without causing fission. Since a critical mass is by definition one in which a neutron has less

than a 35-40% chance of escaping without causing fission, the fusion reaction cannot occur there. Consequently the fusion reaction must take place in a location outside the fissile core.

Simulations show that at the temperatures reached by a 250 ton fission explosion, and at normal densities (gas highly compressed to near liquid density, or in lithium hydrides) even deuterium-tritium fuel does not fuse fast enough for efficient combustion before the expanding fissile mass would cause disassembly. The fuel must be compressed by a factor of 10 or so for the reaction to be sufficiently fast.

Computations also show that care must be taken to heat the fuel symmetrically. The radiation pressure and ablation forces during heating are so large that if significant asymmetry occurs, the fuel will be dispersed before much fusion takes place.

Taken together, these considerations make it evident that neutron bombs are miniaturized variants of staged radiation implosion fusion bombs (see Sections on Thermonuclear Weapons below). The fissile mass is separated from the fusion fuel, which is compressed and heated by the thermal radiation flux from the fissile core. Due to the small mass of the fusion fuel, and the low temperature of ignition, a fission spark plug internal to the fusion capsule is not necessary to ignite the reaction. The ignition probably occurs when the thermal radiation diffuses through the pusher/tamper wall of the fusion capsule. It is also possible that the localized region of intense heating that develops when the shock in the fuel capsule converges at the center may be responsible for, or contribute to, the ignition of the fusion reaction (this is similar to the ignition process in inertial confinement fusion experiments).

The W-79 fissile core is plutonium and is assembled through linear implosion. It is known to contain tungsten and uranium alloys. The likely use of the tungsten is to provide a high-Z material for providing the radiation case, and for the fuel capsule pusher/tamper. Uranium may be used simply to provide inertial mass around the core compression system, it may also serve in part as a neutron reflector.

A notional sketch of the W-79 is given below. The dimensions in centimeters are given along the left hand and lower border of the design. Typical screen formatting will tend to stretch the graphic vertically since line width:character width ratios are usually something like 5:3.

0CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC

9CCCCCCCCCCCCCCCCCCCCCCCCCCCRRRRRRRRC

8CCEEEEEEEEEEEEEEEE RRRR RRRRC

7CCEEEEEEEEEEEEEEEE RRR RRRC

6CCEEEEEEEEEEEEEEEE RRC

5CCEEEEEfffffffEEEEE RRC

4CCEEEfffffffffffEEERRRR RRC

3CCEEfffffffffffffEERRR HH RC

2CCEEfffffffffffffEERR HHHH RC Ogive End ->

1CCEEfffffffffffffEERR HHHHHH RC (pointy end)

0CCEEfffffffffffffEERR HHHHHH RC

9CCEEfffffffffffffEERR HHHH RC

8CCEEfffffffffffffEERRR HH RC

7CCEEEfffffffffffEEERRRR RRC

6CCEEEEfffffffffEEE RRC

5CCEEEEEEEEEEEEEEEE RRC

4CCEEEEEEEEEEEEEEEE RRR RRC

3CCEEEEEEEEEEEEEEEE RRRR RRRRC

2CCCCCCCCCCCCCCCCCCCCCCCCRRRRRRRRRC

1CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC

01234567890123456789012345678901234567

Legend:

C - casing (steel and uranium?)

E - explosive

f - fissile material (plutonium)

R - radiation shield/radiation case (tungsten)

H - hydrogen fuel capsule, made of tungsten, filled with D-T gas

The fissile material mass in this design would be something like 10 kg. The 750 ton fusion yield indicates at least 10 g of D-T mixture for the fusion fuel. Under high static pressure hydrogen

can reach densities of around 0.1 mole/cc (0.25 g/cm^3 for DT). This indicates a fuel capsule volume of at least 40 cm^3, or a spherical radius of 2.5-3 cm including wall thickness.

4.3.3 The Alarm Clock/Layer Cake Design The earliest and most obvious idea for using fusion reactions in weapons is to surround the fission core with a fusion fuel. The radiation dominated shock wave from the expanding fission core would compress the fusion fuel 7- 16 fold, and heat it nearly to the same temperature as the bomb core. In this compressed and heated state a significant amount of fusion fuel might burn.

Calculations quickly showed that only one reaction ignited with sufficient ease to make this useful - the deuterium-tritium reaction. The cost of manufacturing tritium relative to the energy produced from the fusion reaction made this unattractive.

Two ideas were later added to this concept to make a practical weapon design: The first: use lithium-6 deuteride as the fuel. The excess neutrons released by the fission bomb will breed tritium directly in the fuel blanket through the Li-6 + n -> T + He-4 + 4.78 MeV reaction. A layer at least 12 cm thick is necessary to catch most of emitted neutrons. This reaction also helps heat the fuel to fusion temperatures. The capture of all of the neutrons escaping ahead of the shock wave generates about 2.5% as much energy as the entire fission trigger release, all of it deposited directly in the fusion fuel.

The second: encase the fusion fuel blanket in a fusion tamper made of uranium. This tamper helps confine the high temperatures in the fusion blanket. Without this tamper the low-Z fusion fuel, which readily becomes completely ionized and transparent when heated, would not be heated efficiently, and would permit much of the energy of the fission trigger to escape. The opaque fusion tamper absorbs this energy, and radiates it back into the fuel blanket. The high density of the fusion tamper also enhances the compression of the fuel by resisting the expansion and escape of the fusion fuel.

In addition the uranium undergoes fast fission from the fusion neutrons. This fast fission process releases far more energy than the fusion reactions themselves and is essential for making the whole scheme practical.

This idea predates the invention of staged radiation implosion designs, and was apparently invented independently at least three times. In each case the evolution of the design seems to have followed the same general lines. It was first devised by Edward Teller in the United States (who called the design "Alarm Clock"), then by Andrei Sakharov and Vitalii Ginzburg in the Soviet Union (who called it the "Layer Cake"), and finally by the British (inventor unknown). Each of these weapons research programs hit upon this idea before ultimately arriving at the more difficult, but more powerful, staged thermonuclear approach.

There is room for significant variation in how this overall scheme is used however.

One approach is to opt for a "once-through" design. In this scheme the escaping fission neutrons breed tritium, the tritium fuses, and the fusion neutrons fission the fusion tamper, thus completing the process. Since each fission in the trigger releases about one excess neutron (it

produces two and a fraction, but consumes one), which can breed one tritium atom, which fuses and release one fusion neutron, which causes one fast fission, the overall gain is to approximately double the trigger yield (perhaps a bit more).

The gain can be considerably enhanced though (presumably through a thicker lithium deuteride blanket, and a thicker fusion tamper). In this design enough of the secondary neutrons produced by fast fission in the fusion tamper get scattered back into the fusion blanket to breed a second generation of tritium. A coupled fission-fusion-fission chain reaction thus becomes established (or more precisely a fast fission -> tritium breeding -> fusion -> fast fission chain reaction). In a sense, the fusion part of the process acts as a neutron accelerator to permit a fast fission chain reaction to be sustained in the uranium tamper. The process terminates when the fusion tamper has expanded sufficiently to permit too many neutrons to escape.

The advantage of the once-through approach is that a much lighter bomb can be constructed. The disadvantage is that a much larger amount of expensive fissile material is required for a given yield. Yields exceeding a megaton are possible, if a correspondingly large fission trigger is used. This design was developed by the British. The Orange Herald device employed this concept and was tested in Grapple 2 (31 May 1957). A U-235 fission trigger with a yield in the 300 kt range was used, for a total yield of 720 kt - a boost in the order of 2.5-fold. A variant design was apparently deployed for a while in the fifties under the name Violet Club.

The second approach was adopted by the Soviets and proven in the test known as Joe-4 to the West (actually the fifth Soviet test) on 12 August 1953 at Semipalatinsk in Kazakhstan. This resulted in a very massive, but much cheaper bomb since only a small amount of fissile material is required.

Since there is an actual multiplication effect between the fusion reaction and the tamper fast fission, an improved yield can be obtained at reasonable cost by spiking the fusion layer with tritium prior to detonation.

The Joe-4 device used a 40 kt U-235 fission bomb acted as the trigger and produced a total yield of 400 kt for a 10-fold enhancement, although tritium spiking was partly responsible. 15-20% of the energy was released by fusion (60-80 kt), and the balance (280-300 kt) was from U-238 fast fission. A later test without tritium spiking produced only 215 kt.

This design has a maximum achievable yield of perhaps 1 Mt (if that) before becoming prohibitively heavy. The USSR may never have actually deployed any weapons using this design.

4.4 Elements of Thermonuclear Weapon Design

In the previous subsection (4.3) I discussed weapon designs that employ the easy-to-ignite D-T reaction. In principle large fusion explosions could be created using this reaction, if sufficient tritium were available. The fact that tritium must be made through neutron reactions (or other even more expensive charged particle reactions) makes its cost prohibitively high for this. A

neutron expended in breeding Pu-239 or U-233 would make ten times as much energy available for a nuclear explosion. Even if the fusion neutrons were used efficiently in causing U-238 fast fission (requiring a massive fusion tamper as in the Alarm Clock/Layer Cake design), the energy gain would still not be dramatically greater than breeding fissile material directly. For at most a modest energy gain, this design would have considerable penalties. First, there is the added complexity compared to a pure fission bomb. Second, and more important, is the natural decay of tritium. If the weapon is intended to be kept in stockpile rather than used immediately (which fortunately, has been the case since 1945), then maintaining a given tritium inventory means duplicating the initial investment in tritium production every 17.8 years (sic, this is not a typo, the half-life is 12.33 years but continuously replacing decayed tritium requires duplicating production over 12.33/ln 2 years).

4.4.1 Development of Thermonuclear Weapon Concepts

The design and functioning of these weapons is complex, I have decided to introduce the physics and design of these weapons by discussing the evolution of thermonuclear weapon design concepts, instead of plunging directly in to a description of modern hydrogen bomb design principles. This has the advantage of allowing the introduction of important design ideas piecemeal, within a framework which gives the reader a sense of the significance of each.

4.4.1.1 Early Work

In the summer of 1942, quite early in the development of nuclear weapons (3 years before the first fission weapon test), the possibility was noticed of igniting self-supporting thermonuclear combustion in pure deuterium, a naturally available and comparatively cheap material. If the D-D reaction could be initiated then explosions of practically unlimited power could be created inexpensively. An additional possibility was that such weapons might be much lighter than other designs of comparable yield.

Preliminary investigation made the idea seem promising, but more detailed analysis soon showed that the feasibility of a self-sustaining D+D reaction in deuterium at achievable densities was marginal at best. In fact the better part of a decade (until mid 1950) was spent refining calculations to conclusively determine its feasibility one way or the other. In the end, it was shown to be impossible under the conditions then deemed to be achievable.

It was not until early 1951 that a series of conceptual breakthroughs made by Stanislaw Ulam and Edward Teller discovered a way of creating the necessary conditions for solving the "ignition problem". These discoveries led to the detonation of the first hydrogen bomb in November 1952, some 20 months later.

4.4.1.2 The Ignition Problem

Stated as generally as possible, the ignition problem is finding a means of achieving the following requirements.

1. Creating conditions so that the fusion reaction proceeds at a high rate;

2. Maintaining these conditions for a period of time, such that:

3. The total energy produced (rate*time) exceeds the externally supplied

energy consumed in creating and maintaining the reaction conditions.

The formula for the rate of a fusion reaction between two types of particles A and B is:

R = N_A * N_B * f_AB(T)

where R is reactions/sec-cm^3, N_A and N_B are the particle densities of A and B in particles/cm^3, and f_AB(T) is a function the gives the reaction cross section (in cm^2) at temperature T. Since particle energies at equilibrium follow a Maxwellian distribution, the function is actually an average of the cross section values for specific energies over the distribution. In the fusion physics literature, T is normally given in electron volts (eV or KeV). One electron volt of temperature is equal to 11,606 degrees K.

For a given fuel mixture both N_A and N_B are proportional to density. Since the rate at a given temperature is determined by the product of N-A and N_B, it is proportional to the square of the density. A high density can thus greatly enhance the reaction rate.

Note that the effect on the reaction rate in a fixed quantity of fusion fuel only increases linearly with density. This because while the reaction rate (per unit volume) goes up as the square of density, the actual volume decreases linearly with density offsetting this.

The reaction rate is highly dependent on temperature also, as can be seen from the table below:

Reaction Cross Sections (cm^2)

T (KeV)	D/T	D/D	D/He-3
1.0	5.5x10^-21	1.5x10^-22	3 x10^-26
2.0	2.6x10^-19	5.4x10^-21	1.4x10^-23
5.0	1.3x10^-17	1.8x10^-19	6.7x10^-21
6.0	2.6x10^-17	2.3x10^-19	3.3x10^-20
7.0	4.1x10^-17	3.5x10^-19	5.3x10^-20
8.0	6.0x10^-17	5.0x10^-19	8.0x10^-20
9.0	8.2x10^-17	6.7x10^-19	1.3x10^-19
10.0	1.1x10^-16	1.2x10^-18	2.3x10^-19
15.0	2.6x10^-16	1.9x10^-18	1.3x10^-18

20.0 4.2x10^-16 5.2x10^-18 3.8x10^-18

30.0 6.6x10^-16 6.3x10^-18 1.0x10^-17

40.0 7.9x10^-16 1.0x10^-17 2.3x10^-17

50.0 8.7x10^-16 2.1x10^-17 5.4x10^-17

The primary means by which energy is lost from a fully ionized plasma is through bremsstrahlung radiation. The rate of energy emission per unit volume is energy emitted is:

e = 1.42 x 10^-27 Z^3 n_i^2 T^0.5 ergs/cm^3-sec

where n_i is the ion density, Z is the ion atomic number, and T is in degrees K.

Note that the dependence on density for bremsstrahlung emission is the same as that of the fusion reaction rate. Its increase with temperature is comparatively sluggish, varying only with the square root.

4.4.1.3 The Classical Super

The original concept for creating a thermonuclear explosion was to create a thermonuclear combustion wave in a mass of liquid deuterium. The idea was to heat a portion of the mass to ignition conditions with an atomic bomb. The energy released by the burning region would be sufficient to heat an adjacent region to the point of ignition, allowing the region of burning to spread throughout the mass.

This process has been described as "thermonuclear detonation" but "thermonuclear combustion" is more accurate. Unlike a high explosive, the fusion reaction will not go to completion in a narrow zone behind the ignition front, instead the fuel would continue to burn until quenched by the expansion of the fuel mass. Since a large mass takes longer to expand than a small one, it seemed clear that such a bomb must be fairly large for he combustion to be reasonably efficient (one cubic meter of deuterium yielding 10 megatons was the nominal figure used).

Developing this approach (originally called the Super, later the Classical Super) actually required solving two different ignition problems: establishing the initial ignition conditions, and determining whether the combustion wave would be self-supporting once established.

Despite the enormous temperatures and energy densities in a fission bomb, the first problem was not straightforward. Some 80% of the energy in an exploding bomb core is in the form of soft x-rays, but ionized hydrogen at normal densities is virtually transparent to this high temperature radiation (the mean free path is measured in hundreds of meters). The thermal radiation escaping from the bomb core thus cannot heat a localized region of deuterium to ignition. Most the remaining energy is present in the form of kinetic energy of ions and electrons. This energy, transferred as a shock wave by the expanding bomb core, was a possibility.

However the primary means to heat the fuel that was chosen seems to have been the neutron flux from the core, which carried a mere 1% or so of the explosion energy. The reason was that the neutrons travel very rapidly and quickly deposit their energy in a manageably small region (roughly 8 cm thick) through moderation by the hydrogen. Very large explosions were required for sufficient heating to occur (hundreds of kilotons).

The second problem, determining whether thermonuclear combustion would propagate and lead to efficient combustion, was very difficult to solve. A large number of physical processes are involved: questions of energy production rates, energy transport by various reaction products (different types of ions and neutrons of various energies), energy transport by electrons, energy loss from photons through bremsstrahlung and inverse Compton scattering, etc.

Early on it became apparent that direct heating of deuterium could not establish ignition conditions, even fission bombs are not hot enough for this. Adding extremely costly tritium to the ignition zone as a "starter fuel" was required, the easily combustible tritium could in principal raise the fuel to deuterium-deuterium fusion temperatures. Since it is impractical to use tritium as the principal fusion fuel, it was essential that the detonation wave be able to propagate itself into pure deuterium. Determining whether or not this was even possible proved extremely difficult.

Study of the problem began during WWII, and continued until late 1950. The basic problems were the balance of energy production to energy loss, and spread of combustion conditions. In the burning zone, energy emitted as bremsstrahlung (and to a lesser extent the inverse Compton effect) was effectively lost to the fuel mass since it was very unlikely to be reabsorbed (the MFP for these high temperature photons was measured in kilometers!). The fuel mass was thus out of equilibrium with thermal radiation. To make the ratio between production and loss favorable, and the rate of deuterium combustion reasonably rapid, very high temperatures were needed. The energy loss of course tended to damp the temperature rise, making the conditions difficult to achieve and maintain.

The fact that most of the fusion energy was released as neutron kinetic energy was no doubt problematic also. This meant that most of the energy was deposited in a fairly large region outside of the combustion zone, making propagation of the zone more difficult.

Solving these problems required intensive numerical computations to simulate the propagation and combustion process. As computations became more refined during 1949-1950 the ignition problem became worse, requiring ever larger amounts of tritium to reach ignition conditions (eventually it was estimated that 3-5 kilograms were needed - basically filling the entire ignition zone with a 50:50 D-T mixture, corresponding to lost production of 220-500 kg of plutonium). But even this did not solve the propagation and combustion efficiency questions.

Eventually in mid-1950 (after 8 years of study) it became clear that despite large amounts of tritium starter fuel, at best very low combustion efficiencies could be obtained with bombs of

reasonable size. The reaction cross sections of the deuterium-deuterium reactions were simply too low by a factor of 2 to 3 to make a Classical Super feasible.

4.4.1.4 The Teller-Ulam Design

Here the matter rested until January 1951. No viable technical approach for exploding deuterium was available. At this time Stanislaw Ulam was considering ways of improving fission bombs. Since these weapons generally rely on compression, he contemplated whether the energy of a small fission bomb could be used to compress a larger amount of fissile material. Since the energy of fissile material exceeds that of a conventional explosive by six orders of magnitude, if this energy could be harnessed to drive an implosion much more rapid compression and much higher densities could be achieved.

Very quickly he realized that this idea could be extended to compressing deuterium, to make a fusion explosion possible. This is the origin of the key ideas of separation and staging: separating fuels into physically discrete units, and using the explosion of one stage to drive the second stage.

It is not immediately evident from a cursory examination of the physics of thermonuclear detonation waves that this is really of any help. At a constant temperature, the reaction rate in a fusion fuel mass increases linearly with compression, but so does the emission rate of thermal energy through bremsstrahlung. The balance of energy production to energy loss remains the same. For this reason Teller had long regarded compression as being futile for enhancing the classical Super.

However not all physical factors scale similarly. While the rates of energy generation and emission are linear with density, the scale of the whole system varies inversely with the cube root of density just as it does in fission cores. Compressing the fuel by a factor of 1000 (for example), reduces the dimensions by a factor of 10. This has several important consequences. First, the fuel has greater opportunity to burn before disassembly. Second, the MFPs for neutrons decrease by a factor of 1000, and for photons by a factor of a million. Neutron heating thus occurs in a narrower zone, assisting the propagation of the burn region, while photon absorption becomes an important heating mechanism - effectively eliminating bremsstrahlung loss.

The net result is that compression does indeed make a big difference in the feasibility of propagating a thermonuclear combustion wave.

Ulam was the key figure involved with the detailed computations that killed the Classical Super concept, he was thus well positioned to realize the benefits of compression. The fact that he did not propose compression as a solution earlier can be explained by the fact that chemical explosives are too weak to be helpful. Much higher compressions are needed. Once he thought of nuclear driven implosion, the idea of harnessing it to fusion was immediate. He quickly persuaded Teller of the fundamental soundness of the idea.

The next problem was determining how the second stage implosion should actually be carried out. Ulam's concept did not specify how an implosion could be successfully produced using the nuclear explosion energy. His initial idea was to use the kinetic energy in the shock wave of expanding fission trigger debris. Reflecting and concentrating this shock wave on the second stage is possible in principle, but likely to be exceedingly difficult in practice.

Teller soon conceived of a better idea. He had been studying radiation transport in fission explosions and was well acquainted with the physics involved. He knew that most of the energy is in the form of thermal radiation, not kinetic energy, and furthermore that the shock front rapidly emits this energy into the bomb casing. He also recognized that this thermal energy could be harnessed to perform work on the second stage much more easily and efficiently than the kinetic shock.

The idea Teller developed is now known as radiation implosion. The thermal radiation escaping from the primary stage (also called simply the "primary" or "trigger") flows along a gap between the fusion fuel and the opaque bomb casing (known as the radiation channel) until the interior of the casing is heated to a uniform temperature. The blackbody radiation emission from the casing evaporates material from an opaque pusher/tamper around the fusion fuel. The expansion of this heated material acts like a rocket engine turned inside out - the inward directed reaction force drives the fuel capsule inward, imploding it.

Once the idea of separation and staging have been developed, the idea of radiation implosion is actually rather difficult to avoid. The thermal radiation arrives well ahead of the shock, and must be dealt with in some way. It is very unfavorable to allow it to heat the fusion fuel prior to compression, since entropic heating makes compression much less effective. If an opaque radiation shield is placed around the fuel to protect it from heating, the evaporation of the shield and a resulting implosion is inevitable.

A final additional question remains to making this scheme work. How to heat the compressed fuel to ignition temperature? One possibility is achieving sufficient heating from the compression process itself, reminiscent of a diesel engine. Adiabatic compression raises the temperature, but even with extreme compression not by a large enough factor. The extremely rapid implosion necessarily generates an intense convergent shock wave in the fuel. When this shock converges at the center, the extreme heating can be sufficient to ignite the fuel (this approach is used in the radiation imploded fuel capsules used in inertial confinement fusion experiments).

Teller thought of an additional element to the design to accomplish ignition. He proposed placing a sub-critical fissile mass (called the "spark plug") at the center of the fusion fuel. The implosion process would compress this mass to a high level of criticality, causing an extremely rapid fission reaction. This would directly heat the highly compressed fuel, initiating a thermonuclear burn.

As J. Carson Mark points out, the spark plug idea is a fairly obvious addition. After all, it was the idea of compressing fissile material that set Ulam upon this path in the first place, and heating

fusion fuel in direct contact with a fission explosion is the same approach as the original Super concept.

Taken together these ideas form the basis of the "Teller-Ulam" design, more technically described as "staged radiation implosion". So far as is known all high yield nuclear weapons today (>50 kt or so) use this design. It is striking that once Ulam's initial insight regarding the use of a nuclear explosion to compress the fuel was made, the other parts of the concept seem to develop almost inevitably from the effort to translate the concept to practice (which partly explains its reinvention by the Soviets, British, French, and Chinese).

Although the prospect of making possible a self-supporting thermonuclear detonation wave appears to have been the initial attraction to both Ulam and Teller, it turns out that once the final Teller-Ulam concept was developed the character of the ignition problem was so completely changed that this issue ceased to be of major importance.

To permit radiation implosion, and prevent premature heating of the fusion fuel, an opaque tamper is placed around the fuel mass to keep thermal radiation out. By the same token it acts as a radiation container to keep thermal radiation in. Because of this, the issue of balancing energy production and radiation loss is no longer important. The energy produced by the fusion reactions remains trapped inside the tamper, allowing the temperature and reaction rate to rise continuously. This is in fact essential in making the fissile spark plug viable. The very high Z fissile material radiates thermal energy at an extraordinary rate (over seven hundred thousand times faster than hydrogen), and would quench the fusion reaction if the energy could escape.

The ignition problem for the radiation implosion approach now resembles the efficiency problem in fission bombs. The efficiency of the fusion burn is determined by the fusion rate, integrated over the duration of confinement. The fusion process is usually shut down when the fuel capsule undergoes explosive disassembly in a manner similar to that of a fission core. If the reaction is highly efficient it may burn up so of the much fuel that the rate drops off to a negligible value despite the increasing temperature before disassembly occurs.

Unlike the fission bomb though, convenient efficiency equations cannot be analytically derived. The energy release in a fission reaction is governed by a simple exponential function of time. In contrast fusion reactions are not chain reactions, and the way reaction rate varies with temperature is not simple. Further the energy release in deuterium is due to three different reactions, each with a different rate, and the composition of the fuel continuously changes. An adequate treatment of efficiency necessarily relies on numerical simulations.

The fully developed Teller-Ulam design was dubbed the "equilibrium thermonuclear" or "equilibrium super". The meaning of this term is open to question. Some writers (Rhodes in _Dark Sun_ for example) have interpreted it to mean that a dynamical equilibrium is established between the exploding spark plug and the collapsing fuel mass, bringing the fusion fuel to its highest state of compression. This is possible, but I believe the term most likely simply refers to the fact that the burning fusion fuel remains in equilibrium with thermal energy, unlike the Classical Super.

This short overview provides an understanding of the components of thermonuclear weapon designs, and an understanding of the role of each component. It scarcely does justice to the physical processes involved. This requires a more detailed look at each part of the system.

4.4.2 Schematic of a Thermonuclear Device

Below is a representative deptiction of how a Teller-Ulam device is constructed. This schematic is based on a cylindrical design for the secondary. This illustration is probably most representative of the large, high yield designs developed in the early fifties. The most modern and compact designs use spherical secondaries and other geometric variants are possible.

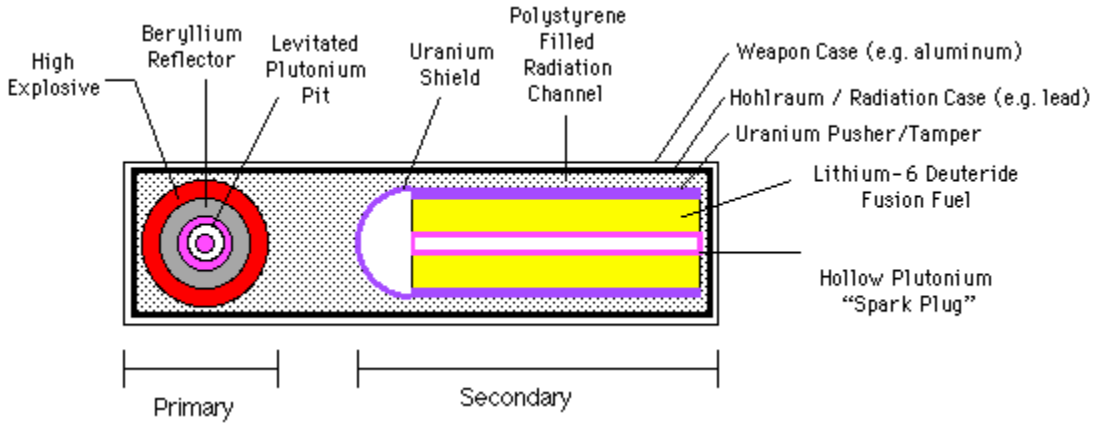

Components of the Teller-Ulam design:

- External Casing (made of structural material: steel, aluminum, plastic, etc.)
- Primary (fission trigger)
- Radiation Shield (high-Z material: uranium or tungsten; this may also contain boron-10 as a neutron absorber)
- Hohlraum or Radiation Case (high-Z material: uranium, lead, or tungsten, etc.)
- Radiation Channel (gap between the casing and the fusion pusher tamper; basically empty, often filled with plastic foam)
- Fusion Pusher/Tamper (high-Z material: natural/depleted uranium, HEU, tungsten, lead, etc.)
- Fusion Fuel (usually Li-6 deuteride; also natural lithium deuteride, liquid deuerium, etc.)
- Spark Plug (fissionable rod of HEU or plutonium)

4.4.3 Radiation Implosion

4.4.3.1 The Role of Radiation

At the temperatures achievable in the fission core of the primary (up to 10^8 degrees K) nearly all of the energy is present as a thermal radiation field (up to 95%) with average photon energies around 10 KeV (moderately energetic X-rays). Most of this thermal energy is rapidly radiated away from the surface of the "X-ray fireball", composed of the expanding X-ray opaque

material of the core and tamper. It is this powerful flux of energy in the form of X-rays that is harnessed to compress the fusion fuel.

To do useful work, the radiant energy from the primary must be kept from escaping from the bomb before the work is completed. This is accomplished by the radiation case - a container made of X-ray opaque (high-Z, or high atomic number) material that encloses both the primary and secondary. The gap between the radiation case and other parts of the bomb (mostly the secondary) is called the radiation channel since thermal radiation travels from to other parts of the bomb through this gap.

The X-ray flux from the primary actually penetrates a short distance into the casing (a few microns) and is absorbed, heating a very thin layer lining the casing to high temperatures and turning it into a plasma. This plasma re-radiates thermal energy, heating other parts of the radiation channel farther from the primary.

The radiant energy emitted by the primary is blackbody radiation: a continuous spectrum of photons whose energy distribution is determined solely by the temperature of the radiating surface. The average photon energy, and the energy of the peak photon intensity are proportional to the temperature. Similarly, the photons re-radiated by the surfaces lining the radiation channel form a blackbody spectrum.

As energy flows down the radiation channel, the energy density drops since the photon gas is now filling a greater volume. This means the temperature of the photon gas, and the average photon energy must drop as well. From an initial average energy of 10 KeV, the X-rays soften to around 1-2 KeV. This corresponds to a temperature in the casing of some 10-25 million degrees K.

[Note: Many descriptions in the open literature exist dating back to the late seventies claiming that energetic X-rays from the primary are absorbed by the radiation casing (or plastic foam), and are re-emitted at a lower energy - implying that some sort of energy down-shifting mechanism (like X- ray fluorescence) is at work. This is a misconception. The lining of the casing is in local thermal equilibrium with the energy flux impinging on it, and re-radiates X-rays with the same spectrum. The X-ray spectrum softens simply because the photon gas cools as it expands to fill the entire radiation channel.]

In physics a closed container of radiation, like the radiation case, is called a "hohlraum". This German word for "cavity" (which has the obvious English cognates "hole" and "room") has been attached to the study of the thermodynamics of radiation since the last century in connection with blackbody radiation. German physicists early in this century used it as a theoretical model for deriving the blackbody radiation laws from quantum mechanics. Energy in a hohlraum necessarily comes into thermal equilibrium and assumes a blackbody spectrum. This is important for obtaining the necessary symmetry for an efficient implosion. Regardless of how uneven the initial energy distribution within the casing is, the radiation field will quickly establish thermal equilibrium throughout the casing - heating all parts to the same temperature.

4.4.3.2 Opacity of Materials in Thermonuclear Design Since the emission, transport, and absorption of thermal radiation is critical to all phases of operation of a thermonuclear device the opacity of various materials to this radiation is critically important. The interaction between a given element and an X-ray photon is dependent on the atomic number, atom density, and ionization state of the element, and the energy of the photon. Since the X-ray flux has a continuous spectrum, we are really interested in the average interaction across that spectrum. We are also especially interested in the situation where the average photon energy (the radiation temperature) and the average kinetic energy of atoms/ions are the same. This situation is called local thermodynamic equilibrium (LTE), and is found almost everywhere inside a thermonuclear device.

The terms "high-Z" and "low-Z" come up frequently in discussing the interaction between thermal radiation and physical materials. These terms are relative - whether a material qualifies as having high or low atomic number depends on the temperature under discussion. The two terms can also be taken as approximate synonyms for "opaque" and "transparent". This is not universally true, however. As explained below, at extremely high pressure this distinction may become unimportant.

An ion strongly interacts with the X-ray spectrum (is opaque to it) when it possesses several electrons, because it then has many possible excitation states, and can absorb and emit photon of many different frequencies. A material where the atomic nuclei are completely stripped of electrons must interact with X-ray photons primarily through the much weaker processes of bremsstrahlung or Thomson scattering. High atomic number atoms hold on to their last few electrons very strongly (the ionization energy of the last electron is proportional to Z^2), resisting both thermal ionization and high pressure dissociation, which is the primary reason they are opaque.

Even when comparing two different fully ionized materials, the higher Z material will more readily absorb photons since bremsstrahlung absorption is proportional to Z^2 (at equal particle density, if it is the ion densities that are the same then it is proportional to Z^3). See Sections 3.2.5 Matter At High Temperatures, and 3.3 Interaction of Radiation and Matter for more discussion of these are related issues.

To summarize: a material qualifies as opaque or "high-Z" if it possesses some electrons at the temperature under consideration. A transparent or "low-Z" material will be completely ionized. Since electrons are removed by ion/particle collisions, the ionization state will depend on the temperature, which is determined by the average kinetic energy (kT) of a particle. At a minimum, all electrons with ionization energies less than or equal kT will be removed, and at the densities of matter encountered here electrons with ionization energies up 3 or 4 kT will often be removed as well.

An important caveat to the above is that at the very high pressures that exist in a fully compressed secondary, essentially any element will become opaque. The density of Fermi degenerate matter under some specific pressure is determined by the density of free electrons. Under the enormous compressive forces generated during secondary implosion the electron

density becomes so high that even the "weak" Thomson scattering effect becomes strong enough to render matter opaque. This is important for the energy confinement needed during the thermonuclear burn.

Different temperatures are encountered in different parts of a thermonuclear device - approaching 10 KeV in the primary, 1-2.5 KeV in the radiation channel, and up to 35 KeV inside the secondary. We can make a general guide showing which materials qualify as opaque or transparent at these temperatures by finding Z such that the last (Zth) ionization state has an ionization energy I(Z) approximately equal to kT and 4kT. Any element with I(Z) of around kT will certainly be completely ionized at temperature T. An element will need I(Z) to be significantly greater than 4kT to be highly opaque.

It should be noted that a significant proportion of Planck spectrum energy to be carried by photons with energies even higher than 4kT (10% of it is carried by photons with energies above 6.55 kT). It is possible then at temperatures where the radiation field dominates for the flux of photons to be so intense that photo-ionization of electrons with energies well above 7 kT may occur.

Temperature	Low Z (I(Z)~kT)			High Z (I(Z)~4kT)		
	Z	Symbol	Actual I(Z)	Z	Symbol	Actual I(Z)
1 KeV	9	F	1.10 KeV	18	Ar	4.41 KeV
2.5 KeV	13	Al	2.30 KeV	28	Ni	10.7 KeV
10 KeV	28	Ni	10.7 KeV	55	Cs	41.1 KeV
35 KeV	51	Sb	35.4 KeV	101	Md	139. KeV

From the table we can make some general statements about the materials we want in different parts of the device. We want thermal radiation to escape rapidly from the primary, so it is important to keep the atomic number of materials present in the explosive layer to no higher that Z=28. The use of baratol (containing barium with Z=56) is thus very undesirable. Since the radiation channel needs to be transparent, keeping materials with Z above 9- 13 out of the channel is desirable. Radiation case linings should have Z significantly higher than 55, as should the fusion tamper and radiation shield.

Due to the complexity of the interacting processes that determine the opacity of incompletely ionized material at LTE, theoretical prediction of these properties is extremely difficult. In fact accurate predictions based on first principles is impossible, experimental study is required. It is interesting to note that opacity data for elements with Z > 71 remain classified in the US. This is a clear indication of the materials used in thermonuclear weapon design for containing and directing radiation. The fact that elements with Z > 71 are used as radiation case linings has recently been declassified in the US.

There are 14 plausible elements with atomic number of 72-92 that may be used for this purpose. Of these 14 elements, 5 are definitely known to have been used in radiation case or secondary pusher/tamper designs in actual nuclear devices: tungsten (74), gold (79), lead (82), bismuth (83), and uranium (92). There is evidence that rhenium (75) and thorium (90) may have been used as well, and tantalum (73) has been used in ICF pusher designs. Two others, mercury (80) and thallium (81) are also known to have been incorporated in thermonuclear weapons in classified uses (in addition to declassified uses, such as electrical switches).

The optimal material for radiation confinement should have maximum optical thickness per unit mass. Opacity increases with atomic number, but for a given radiation temperature the increase with Z probably declines at some point. Since atomic mass also increases with Z, there is probably an optimal element for any given radiation temperature that has a maximum opacity per unit mass.

4.4.3.3 The Ablation Process The thin hot plasma layer lining the radiation channel not only radiates heat back into the channel, it also radiates heat deeper into the material lining the channel creating a flow of thermal radiation into the radiation case and the secondary pusher/tamper. The hot plasma also has tremendous internal kinetic pressure and expands into the radiation channel.

This rapid evaporation and expansion (ablation) of the radiation channel lining is unavoidable. Due to the conservation of momentum, the expanding material creates a reaction force called "ablation pressure" that pushes in the opposite direction - blowing the walls of the radiation case outward, and the pusher/tamper of the secondary inward. It is this inward force, analogous to the force exerted by the exhaust of a rocket, that compresses the secondary.

We can calculate representative parameters for the implosion process. To span a range of designs and parameter values let us consider the Mike device, a high yield design that was the first (and undoubtedly physically largest) radiation implosion device ever exploded, and the W-80 cruise missile warhead which is a modern light weight design.

The casing of the Mike device was a steel cylinder 20 ft. (6.1 m) long and 80 in. (2.0 m) wide, with walls 12 in. (30 cm) thick. It used a TX-5 fission primary, with a yield probably no larger than 50 kt, and produced a total yield of 10.4 Mt. The W-80 is a cylinder 80 cm long, and 30 cm wide, it has a primary with a yield in the low kiloton range (call it 5 kt for the sake of the discussion), and a total yield of 150 kt. The thickness of the W-80 casing is unknown, but given its weight (130 kg) it must be less than 2 cm.

Once equilibrium is established, the energy density in the radiation channel will be roughly the energy released by the primary, divided by the volume inside the radiation case (this neglects the kinetic energy in the primary remnants, and the volume of the secondary, but these are comparatively small and offset each other). This gives radiation densities of 2.2 x 10^14 erg/cm^3 for Mike and 4.3 x 10^15 erg/cm^3 for the W-80, a energy density ratio of 1:20. By applying the blackbody radiation laws (see Section 3.1.6 Properties of Blackbody Radiation) we can determine the corresponding radiation intensities and temperatures: 9.8 x 10^6 K and 5.3 x

10^16 W/cm^2 for Mike; and 2 x 10^7 K and 1.0 x 10^18 W/cm^2 for the W-80. The radiation pressures are 73 and 1400 megabars respectively.

The ablation pressure is determined by mass evaporation rate, and the effective exhaust velocity of the evaporated material:

P = m_evap_rate * V_ex

If the evaporation rate is in g/cm^2-sec, and V_ex is in cm/sec, then the result is in dynes/cm^2, applying the conversion factor of 10^6 dynes/cm^2 per bar gives the result in bars.

The ultimate implosion velocity is determined by the rocket equation:

V_imp = V_ex * ln(m_initial/m_final)

where m_initial is the initial mass of the pusher/tamper and m_final is the mass after ablation is complete. Peak efficiency (in terms of energy expended) of an ideal rocket is reached when the ratio (m_initial/m_final) is around 5.

In a rocket maximum force is extracted from hot reaction gases by allowing them to expand as they exit the rocket nozzle, which cools and accelerates the exhaust. The effective exhaust velocity is the velocity of the cooled and expanded gas at the nozzle's mouth. In contrast, ablation is generated by an energy flow that must penetrate the exhaust gas which prevents the gas from cooling. The effective exhaust velocity here is the gas velocity at the sonic point, the point where the gas is moving at the local speed of sound relative to the ablation front, where the material is actually evaporating. Since changes to the exhaust flow beyond the sonic point cannot propagate back to the ablation front, as far as the secondary is concerned the exhaust effectively disappears at this point.

[Note that many descriptions in the open literature ascribe the driving force in implosion to the plasma pressure created by a plastic foam that is known to fill the radiation channel in some weapon designs. Since hydrodynamic effects that occur beyond the sonic point cannot propagate back to the imploding secondary, this is impossible.]

Since the exhaust gases beyond the sonic point absorb heat and carry it away from the secondary, and also reradiate significant amounts of thermal energy back into the radiation channel, the ablation driven acceleration process is less efficient than an ideal rocket as judged in terms of the incident radiation intensity.

The efficiency for an ideal rocket (the percentage of the kinetic energy in the exhaust-rocket system ending up in the rocket at burnout) is given by:

eff = (x (ln x)^2)/(1 - x)

where x if the ratio between the final mass and the initial mass:

x = m_final/m_initial

This has a peak efficiency of 64.8% at x = 0.203.

The heating of the exhaust limits the ablation driven rocket to a maximum efficiency of approximately 15-20% when x is in the range of 0.1 to 0.6 (with peak efficiency around 0.25). Above 0.6 it drops off to about 7% at 0.85. It thus desirable to ablate off most of the pusher/tamper mass so that x < 0.5. [Note: This is based on ICF data which uses radiation driven implosions at a few hundred eV. The higher temperature X-rays of nuclear implosion systems penetrate to the ablation front more efficiently and may actually do better than this.]

Scaling laws for the relationships between temperature or energy density and the ablation rate and exhaust velocity can be determined by dimensional analysis. The sonic-point temperature (and average kinetic energy) is proportional to (~=) the temperature in the radiation channel, and since

v ~= KE^(1/2),

then

V_ex ~= T^(1/2).

Because the incident energy flux I between the ablation front and the sonic point must be proportional to the kinetic energy carried away we have:

I ~= m_evap_rate * V_ex^2 ~= m_evap_rate * T

and since

I ~= T^4

we get

m_evap_rate ~= T^3.

Finally:

P = m_evap_rate * V_ex ~= T^3 * T^(1/2) ~= T^3.5

It is possible to estimate the values of the constants to convert these proportionalities into equations from physical data, but the process is rather elaborate. We can borrow some relationships that have appeared in the inertial confinement fusion literature in connection with radiation implosion to get some estimates of the magnitudes:

P (bars) = 0.3 T^3.5

and

m_evap_rate (g/cm^2-sec) = 0.3 T^3

where T is in electron volts. From this we get:

V_ex (cm/sec) = P/m_evap_rate = 0.3 T^3.5 (10^6 dynes/cm^2 /bar)/0.3 T^3

= 10^6 T^0.5

For the Mike device this gives:

P = 5.3 x 10^9 bars

m_evap_rate = 0.18 g/cm^2-nanosecond

V_ex = 2.9 x 10^7 cm/sec = 290 km/sec

For the W-80:

P = 6.4 x 10^10 bars

m_evap_rate = 1.5 g/cm^2-nanosecond

V_ex = 4.1 x 10^7 cm/sec = 410 km/sec

The ablation pressures for the Mike and W-80 devices are much greater than the corresponding radiation pressures, by factors of 73 and 46 respectively. This shows that the force exerted by radiation pressure is comparatively small.

From the classical rocket equation given above we can estimate V_imp at maximum efficiency (where 75% of the mass is ablated off) at 400 km/sec (Mike) and 570 km/sec (W-80).

4.4.3.3.1 The Ablation Shock

There is a short "settling" period early in the implosion process when the initial ablation pressures are propagating through the pusher/tamper. When the radiation flux begins ablating the pusher, a strong shock wave propagates through the pusher/tamper. This shock compresses and accelerates the tamper inward.

When the shock reaches the inner surface of the tamper, the tamper is (more or less) uniformly compressed and at its minimum thickness. The material beyond the tamper has a much lower density, so the shock compressed material, which is under extremely high pressure, immediately begins to expand and form a release wave (see Section 3.6.1.1 Release Waves). This release wave has two edges, a forward edge where the expanding gas meets the low density material, and a rear edge where the pressure drop begins in the shock compressed gas.

At the forward edge of this wave most of the internal energy of the gas has been converted to kinetic energy. This means that the gas velocity is at a maximum and the pressure has dropped to a minimum. The actual pressure and velocity of the forward edge depends on the density of the low density material. If the density is greater than zero, then this forward edge will be a low pressure shock front.

Behind this leading edge gas velocity decreases, and pressure and density both increase. These changes are continuous, increasing in magnitude with distance from the leading edge until the original state of the shocked gas is reached at the rear edge of the release wave. Eventually this rear edge (which travels at the local speed of sound) will reach the ablation front at which point the pressure, density, and velocity distribution in the tamper reaches its final overall form, with a continuously decreasing pressure and density gradient from the ablation front forward to the leading edge of the tamper.

This pressure gradient is responsible for the inward acceleration of all of the material that has passed through the inner implosion shock front. The general pressure and density profile, once established, remains stable throughout the implosion process, until the inner front collides with itself at the center.

It is likely that while the initial shock is moving through the tamper, the pressure at the ablation front will continue to climb, creating a pressure and density gradient behind the shock. The pressure gradient created by the release wave will merge with this compression wave to create the continuous pressure and density gradient.

The velocities of the shock waves generated in a uranium pusher/tamper by these pressures is on the order of 150 km/sec (for Mike) and 550 km/sec for the W-80. These shocks are powerful enough to dissociate electrons from their nuclei by pressure alone, but they also strongly heat the tamper (causing thermal ionization) and prevent the achievement of true Fermi degenerate compression (see below). Still, with the effects of ionization, the density increase can be greater than a factor of 10. If the rise in ablation pressure can be moderated by appropriate design techniques (see Implosion System Design below) so that the shock front can traverse the tamper ahead of the full pressure jump, this can be substantially improved upon. The temperatures produced are in the order of several million degrees K, considerably lower than the radiation channel temperature. The pressure behind the shock front is predominantly due to particle kinetic pressure not radiation pressure, so these are simple supercritical radiative shocks.

4.4.3.4 Principles of Compression

4.4.3.4.1 Purpose of Compression The fundamental purpose of compressing the fusion fuel is to allow the reaction to proceed swiftly enough for a large part of the fuel to burn before it disassembles from its own expansion, or from the expansion of the trigger fireball. Compression promotes the fusion reaction in several ways.

In Section 4.4.1.2 (The Ignition Problem) it was seen that the reaction rate of a quantity of fuel will increase in direct proportion to its density if it is compressed at constant temperature.

But we can expect the temperature to increase as well, and with it the cross-sections of the thermonuclear reactions. In dense thermonuclear fuel nearly all of the energy present exists as a photon gas. Since the radiation energy density in a sealed container is dependent only on temperature, confining the energy to a smaller volume increases the temperature. According to the blackbody radiation law (Section 3.1.6 Properties of Blackbody Radiation) this increase is rather slow though, being proportional to the fourth root of the energy density. This is offset by the fact that throughout most of the temperature range of interest the increase in cross section with temperature is rapid.

To illustrate the relative importance of these effects, suppose a quantity of fuel were compressed 16-fold. This would increase the reaction rate 16 times due to the density increase alone. This compression would double the temperature (16^(1/4) = 2). If the temperature were initially 5 KeV (58 million degrees K), the cross-section increase for the pace-setting D-D reaction would be 6.7 fold (although this factor would decline as the temperature rises).

These are not the only advantages however. Just as it does with a fissile core, compression of the fusion fuel increases the dimension of the fuel mass as measured by neutron collision mean free paths. The neutrons released by the fusion reactions will thus undergo many more collisions with fuel nuclei before they can reach the tamper. In the early stages of combustion of pure deuterium fuel (before the temperature rise and buildup of He-3 make the He-3 + D reaction significant) 66.3% of all the energy produced is released in the form of neutron kinetic energy. Deuterium is a very light atom (only hydrogen-1 is lighter) so it has a very strong moderating effect. On average a neutron will lose 51.6% of its energy with each collision (see Section 4.1.7.3.2.1 Moderation and Inelastic Scattering). After several collisions then, almost all the energy released as neutron kinetic energy will be transferred to heating the fusion fuel.

The MFP in liquid deuterium for the 14.1 MeV neutrons produced by the D+T reaction is 22 cm. A 1 kg sphere of liquid deuterium would be 22.4 cm across, most 14.1 MeV neutrons generated within this mass would escape without even a single collision. If this sphere were compressed 125-fold, its diameter would shrink to 4.49 cm but the MFP would now be only 0.18 cm. Few neutrons would escape without depositing most of their energy in the fuel mass.

The effects of multiple neutron collisions are even more important in lithium deuteride fuel. 75% of the mass of Li-6 D fuel consists of lithium. To use this mass as fuel, the lithium nuclei must each capture a neutron. The fuel must be compressed sufficiently for a large part of the neutrons produced to participate in this reaction.

To illustrate how compression affects the rate of burn up, I have run a simple computer model of deuterium fusion at varying densities. The model assumes constant density during the fusion reaction and no energy escape from the fuel.

With the same amount of energy deposited in the fuel for ignition (0.1 kilotons/kg) the time to burn up 75% of the deuterium at normal liquid density (0.16 g/cm^3) is 1.3 milliseconds. At 288 g/cm^3 the time shrinks to only 4.4 nanoseconds. The graph of density versus time is nearly a straight line on a log-log plot, so intermediate values can be easily estimated using a scientific calculator (or log-log graph paper).

The reduction in burn up time is partly due to the higher initial temperature of the denser fuel (12 million degrees K at 0.16 g/cm^3, and 55 million degrees K at 288 g/cm^3) but even at constant initial temperature the comparative burn up rates are much the same.

4.4.3.4.2 The Fermi Pressure

From the ablation pressures calculated earlier, we can determine the maximum densities that can be produced. Maximum density is achieved if the heating during compression is negligible, that is, the counterbalancing pressure in the compressed material is simply the Fermi pressure produced by a completely degenerate Fermi gas (see 3.2.4 Matter At High Pressures).

Since:

P_Fermi (bars) = 2.34 x 10^-33 * n^(5/3)

where n is the electron density (electrons/cm^3), we can calculate the electron densities of 2.6 x 10^25 electrons/cm^3 for 5.3 gigabars (the Mike device) and 1.2 x 10^26 electrons/cm^3 for 64 gigabars (the W-80). The limiting mass density based on the calculated ablation pressure for the Mike device is thus 86 g/cm^3 (deuterium or Li6D) and 290 g/cm^3 (U-238 at 38% dissociation). The corresponding values for the W-80 are 380 g/cm^3 (deuterium or Li6D) and 1200 g/cm^3 (U-238 at 41% dissociation). The inevitable shock induced entropy increase in the tamper will reduce the achievable densities of a U-238 tamper to values well below this.

Now the energy densities due to degeneracy pressure implied by these electron densities is given by:

E_density (erg/cm^3) = 3.50 x 10^-27 n^(5/3)

or 8.0 x 10^15 erg/cm^3 for 5.3 gigabars, and 1.0 x 10^17 erg/cm^3 for 64 gigabars. The Fermi energy per mass at 5.3 gigabars is 9.3 x 10^9 J/kg (D or Li6D) and <7.3 x 10^9 J/kg (U-238), or 2.2 and <1.7 tonnes of explosive energy per kg respectively. At 64 gigabars the energy per unit mass is 2.6 x 10^10 J/kg (D or Li6D) and <2.0 x 10^10 J/kg (U-238), or 6.2 and <4.8 tonnes per kg. These figures show the minimum energy investment required to achieve the maximum density.

The corresponding Fermi temperatures, given by:

T_Fermi = (5 P_Fermi)/(2 nk)

are 3.7 x 10^6 K (5.3 gigabars), and 9.7 x 10^6 K (64 gigabars). In efficient (i.e. Fermi degenerate) compression the final temperatures of the compressed fuel most be substantially lower than the Fermi temperature.

4.4.3.4.3 Efficient Compression

Now the question arises as to how these tremendous pressure can be applied to actually generate densities close to these maximum values. Simply applying these ablation pressures suddenly to the thermonuclear fuel will not actually compress it very much. Sudden pressure jumps produce intense shock waves that expend shock energy about equally between heat and kinetic energy, with a negligible portion going to compression. The density increase will limited by the effective value of gamma. Such a violent shock would be radiation dominated so no more than a 7-fold compression occurs in this case.

There are two ways this can be done. The pressure increment can divided into a series of shock waves, each providing a modest pressure increase ratio, and minimal entropic heating. Alternatively, an appropriately shaped continuous pressure rise can produce true adiabatic compression. Actually there is not much difference between these two options. Continuous adiabatic compression is the limiting case of an infinitely large number of infinitely weak shocks producing no entropic heating. And in practice, any continuous pressure gradient of this magnitude will tend to break up into a sequence of discrete shocks (see 3.7.5 Methods for Extreme Compression for further discussion of this).

In any case, propagating an initial shock of considerable strength into the fuel is unavoidable. This is because the bomb will typically disassemble on a time scale of no more than a microsecond or so, and the compression of the fuel must be complete well before this occurs. Even if the fuel layer to be compressed is only a few centimeters thick, then an initial shock of at least some tens of kilometers per second is necessary to traverse the fuel.

We can make some useful observations about the compression of the secondary rather easily.

If A denotes the area of the fuel capsule surface, then the total force being exerted on the capsule is:

F = P*A

From the Newtonian law:

work = force*distance

we can determine the work done on the capsule by:

W = P*A*d = P * (change in volume)

for small values of d. More generally we can say:

W = Integral[P] dV

Now we consider two possible implosion geometries - cylindrical and spherical (variants of these, such as tapered cylinders and ellipsoids are possible but the principles are the same). Since V is proportional to r^2 (for cylindrical geometry) or r^3 (for spherical geometry), it follows that most of the work done on the capsule occurs during the early stages of implosion when the net change in radius is fairly small.

By the time the capsule volume has been reduced by half, then half of all the work that is done on the capsule has been completed (assuming constant ablation pressure). This corresponds to a radius reduction (measured from the outside surface of the pusher/tamper) of:

1 - 0.5^(1/2) = 29.3% (cylindrical geometry)

or

or 1 - 0.5^(1/3) = 20.6% (spherical geometry).

Another way to look at it is that as the capsule implodes, its surface area shrinks. Since the pressure (which is the force per unit area) is constant, the total force, and the ability to do work, on the capsule shrinks also with decreasing radius.

At this point the imploding capsule has acquired half of its final kinetic energy, and 70% (0.5^0.5) of its final implosion velocity. The remaining part of the implosion can be termed "free fall", during which the pusher/tamper travels inward at essentially constant velocity. This maximum velocity depends on the ablation pressure, mass of the pusher/tamper, and the volume of secondary, as well as the geometry.

We should observe that the ablation process soaks up a lot of energy. A simplistic computation of the work done in imploding a 10 liter secondary in the "W-80", assuming constant ablation pressure, shows that 6.4 x 10^13 J (15 kt) is put into the secondary. This is more energy than we assumed that the primary actually produced (5 kt), and also ignores the limits of rocket efficiency. Clearly as implosion proceeds the available energy in the channel decreases, as does the ablation pressure. We can expect the final implosion velocity to be in the range of 300-800 km/sec however.

By the same geometric argument used above, if the compression wave propagates only 30% of the way into a cylindrical fuel mass, or 20% of the way into a spherical one, then half the fusion fuel will have been compressed. Thus the compression process only needs to be efficient in the outer layers of fuel. Even if excessive heating and poor compression occur at significantly smaller radii, overall fuel compression will still be efficient.

From the equations given in Methods of Extreme Compression, and assuming an initial 50 megabar shock though the outer layers of fuel which compresses liquid deuterium or Li6D fuel 12-fold (a 200 km/sec shock in deuterium, a 90 km/sec shock in LiD), we can estimate final compressions and densities for our two illustration cases. The compression for Mike (5.4

gigabars and liquid deuterium) is 197-fold (33.3 g/cm^3); for the W-80 (64 gigabars and Li6D) it is 878 fold (720 g/cm^3). The calculated compression for Mike is well short of the limit imposed by the Fermi pressure, with a lower initial shock pressure than the one I assumed the density could be increased by a factor of two or more. The density calculations for the W-80 (which were done without considering degeneracy effects) is higher than the Fermi pressure. We can conclude then that this system can achieve compressions near the Fermi pressure limited density. [Note that to more properly calculate the effects of compression on dense tamper material at pressures lying between those achieved by high explosives, and the Fermi degenerate state requires use of more complicated theoretical models like the Thomas- Fermi theory.]

4.4.3.5 Ignition

Efficient compression can raise the temperature of the fuel to a few million degrees K. This is hot enough to create a measurable D-D fusion reaction in the compressed fuel, but by itself it does not result in a thermonuclear reaction that is rapid enough to be useful.

To achieve efficient fuel burn up the fuel most be heated to the point where the rate of self-heating becomes significant, triggering a rapidly accelerating combustion process. The denser the fuel mass, the less energy is required to reach this point.

How hot the fuel must be is determined by the density, and the achievable confinement time of the fuel - which in turn is governed by a number of weapon design factors, including the size of the secondary. Using the same simple deuterium fusion model mentioned above, we find an effective ignition temperature of 30 million degrees K. At this temperature the reaction rate and fuel temperature immediately begin a rapid rise, causing accelerating fuel burn-up. The reaction is essentially complete (80% burnup) in 20 nanoseconds when the fuel density is 100 g/cm^3. Lower temperatures create a latency period where the temperature rises very slowly, before abruptly climbing upwards (once 30 million K is reached). At an initial temperature of 12 million degrees, this latency period is 60 nanoseconds after which the fuel burns to near completion in the same 20 nanosecond period. Investigating other densities in the range of 50-300 g/cm^3 gives much the same picture regarding the ignition temperature, although the density does strongly affect how long the fuel burn up takes.

In any case, it is clear that the temperatures prevailing in the ablation- induced shock are much too low to ignite efficient fuel burning.

The energy required to heat the fuel to 3 x 10^7 K is in the range of 2.8 to 4.1 x 10^11 J/kg (67 to 98 tonnes of explosive energy) for deuterium and Li6D fuel with densities between 50 and 200 g/cm^3. This is a factor of 5 times (200 g/cm^3, Li6D fuel) to 15 times (50g/cm^3, D fuel) higher than the energy in the fuel due to degeneracy pressure. Heating the fuel to ignition is thus energetically more expensive than efficient compression.

At least two different mechanisms are possible for igniting the main fusion reaction:

- The first method to be used was suggested by Teller - the use of a fission spark plug in the center of the fuel mass.
- A second method is used in laboratory scale fusion explosions (inertially confined fusion experiments, that is) - allowing the compression shock to converge in the center of the fuel, creating extremely high temperatures and in a very small volume of fuel. A combustion wave then spreads from the center to the remaining fuel. A variant on this to place a tritium-deuterium spark plug in the center of the secondary. Since this reaction ignites at much lower temperatures, it is much easier to achieve the necessary ignition conditions. In ICF experiments D-T mixtures are the only fuel used.

4.4.3.5.1 Fission Spark Plugs

A subcritical fissile mass placed in the center of the fuel will be rapidly compressed upon arrival of the imploding shock. At such a small radius (a few centimeters), the pressure gradient or shock sequence will probably have merged into a single extremely energetic shock. This shock will have been further augmented by the effects of shock convergence, and the final stages of implosion - where the compressed fuel mass decelerates the high velocity tamper - may have generated pressures even higher than the ablation pressure. This shock will have velocities in the range of several hundreds of kilometers a second. When this shock arrives at the interface between the fusion fuel and the higher density spark plug, it will drive an even higher pressure (but slower) shock into the fissile material. The implosion velocity achieved will be at least 100 km/sec in any case - more than an order of magnitude higher than the highest velocities achieved by practical high explosive implosion systems.

In principal a shock of this intensity could compress the spark plug to a density of perhaps 16 times normal, but here too the effects of predetonation intervene to prevent this from being reached. If neutrons are present at the outset of supercritical insertion, the energy release from fission will halt the implosion well short of this density. But - here the much higher velocity of implosion makes a *big* difference in the practical effect of predetonation.

First the enormous kinetic energy and pressures in the imploding mass requires energy releases in the order of a few kilotons simply to halt the implosion process, unlike the high explosive case where the energy release required is negligible compared to the final yield. Second, the compression that is achieved at this point, while much lower than the maximum that the shock is capable of producing, is still probably at least a factor of 3.5 to 4 - as good as that achieved by the best conventional implosion systems under optimum conditions. The result is that an efficient fission explosion should always result.

This is important because neutrons are inevitably present in abundance. First, even though the production of energy in the compressed fusion fuel is negligible at this point, its production of neutrons should be quite significant. Second, and even more important, are neutrons from the trigger explosion. These are leftover stragglers from the primary explosion, now long past (300 nanoseconds or more, such are the time scales with which we are dealing). All of the excess neutrons from the primary (on the order of 10^24) have long since escaped from the expanding

primary fireball, and if nothing has slowed them down they are now well outside the bomb casing.

A substantial fraction of them have, however, have entered the fuel capsule. Unless an absorber has been intentionally placed between the capsule and the primary this neutron population would be on the order of 10^22. The average scattering mean free path for fission spectrum neutrons in liquid deuterium is 7.8 cm, and 4.2 cm in lithium-6 deuteride, so once a neutron enters it will usually scatter repeatedly. Both fuels are very good moderators. With each collision a neutron is robbed of nearly half its energy, on average. Now after a collision, a neutron may escape the capsule or be absorbed by the lithium (the chances of capture by deuterium is negligible) so the population of neutrons declines with time. But since they are losing energy, the time scale for absorption and escape keeps getting longer and longer. The time for complete thermalization is several microseconds, but in the time available before the spark plug fires the neutron energies would still be in the KeV range, and the number of collisions that would have occurred would number scarcely more than a dozen. Even if half of the neutrons were lost after each collision (a high estimate), the neutron population in the fuel capsule would still be astronomically high (>10^17). Since a neutron absorber would have to be implausibly thick (the order of 40 mean free paths) to reduce this to a negligible level, we can assume that many fission neutrons will remain present.

Energy is transmitted from the spark plug to the fuel by both neutrons and photons. The neutron MFP in the Mike model is reduced to 7.8 cm/197 = 0.040 cm, and to 0.0048 cm for the W-80, thus allowing strong neutron mediated heating of the fuel in a thin layer around the spark plug.

From the electron densities calculated above, we can compute the mean free path for Thomson scattering in the compressed secondary at 0.058 cm (5.4 gigabars) and 0.013 cm (64 gigabars). These values are much smaller than the radius of the compressed spark plug, or the thickness of the fuel or tamper layers. The can see that the entire secondary is opaque, strongly scattering the emitted photons and causing photon transport to occur by diffusion. Thomson scattering by itself does not cause fuel heating, the photons must be absorbed before this can occur. The due to the high densities, the spectrum averaged MFPs for the photons flux from the spark plug is quite short also. Assuming a nominal fuel temperature of 10^6 K from compression heating, for the lower compression, low-Z deuterium fuel in Mike we can estimated an absorption MFP of 0.3 cm. For the lithium-containing W-80 fuel it is less than 0.001 cm.

The energy produced by fission will thus be transmitted through the fuel by means of a radiation dominated shock or pure Marshak wave. The fuel will ignite ahead of the full shock heating zone by the leading thermal diffusion zone. Although a spark plug can easily be designed to directly supply sufficient energy to ignite the entire fuel mass, the fact that the heating travels outward by a Marshak wave may allow much smaller spark plugs since the ignition wave may be self-sustaining. The emission of fusion neutrons ahead of the ignition zone may also play a significant role in the growth of the ignition region.

The use of fission spark plugs appears to be the most common (if not exclusive) means of igniting secondaries in deployed designs.

4.4.3.5.2 Shock Heating Induced Ignition

It was noted in Section 3.7.3 Convergent Shock Waves that considerable heating occurs near the center of an implosion, where the shock wave converges in principal to a mathematical point or line. In principal, the temperatures reached at the very center are unlimited.

It is possible to ignite a small mass of fusion fuel in this central region of strong heating. The fusion reactions occurring here can then spread outward through the entire fuel mass as a thermonuclear combustion wave as described above for spark plugs.

It is not clear from present evidence whether this approach has ever actually been used in a real design - either deployed or merely tested. If so, it is likely that a deuterium-tritium mixture would be deliberately introduced at the center to provide a "match" to more easily ignite the fuel. It is known that lithium tritide has been used by the U.S. in thermonuclear secondaries. Since tritium is far too expensive to use in a weapon unless its energy yield is greatly magnified in some way (similar to its use in fusion boosted fission bombs), this may be evidence of the use of this type of ignition system.

4.4.3.6 Burn and Disassembly

Once ignited, the thermonuclear reaction is self-heating. Since the reaction rate increases with temperature, feed-back is established that causes the power output of the secondary to rises steeply. When more than half of the fuel has burned the temperature cannot rise much more since most of the energy has already been released. The depletion of fuel then catches up, and the power output levels off, then begins a somewhat less rapid decline. The period during which the majority of the fuel is burned amounts to a mere 20 nanoseconds or so. All this assumes of course that the disassembly of the secondary hasn't yet intervened to quench the reaction.

At fuel densities on the order of 100 g/cm^3, the maximum temperature can rise to about 350 million degrees K. Under these conditions the pressure tops 100 terabars (100 million megabars, 10^14 bars, or 100 trillion atmospheres). This tremendous temperature and pressure is initially confined to the fusion fuel. It propagates into the tamper as a Marshak wave (a radiation driven compression wave), compressing and accelerating the tamper material outward.

Pressures of this magnitude are capable of generating an outgoing compression wave in the tamper with a velocity of several millimeters per nanosecond. This rate of expansion is so fast that even during the extremely short period when the thermonuclear reaction is near its peak, the density of the fuel could drop significantly and impair overall efficiency. It helps considerably if the tamper is still imploding rapidly when the reaction ignites, since pressure of a few terabars will be necessary to simply to bring the implosion to a halt.

The picture is much the same even in secondary designs where most of the energy is released by fast fission of the surrounding tamper. The pressure in the fusion fuel should be considerably higher than in the tamper, since the fuel energy density of the fissile tamper is substantially lower and it lags behind the fusion reaction slightly (due to the finite velocity of the escaping

neutrons). The pressure in a fissioning tamper can have a substantial confining effect however. The tamper itself will start expanding outward into the radiation channel as it fissions in manner very similar to a disassembling fission bomb core.

Note that the use of fast fission to produce energy in a bomb involves the tamper surrounding the fuel, *not* the bomb casing as is sometimes reported. The highly compressed imploded tamper has an extremely high mass density per unit area and is almost inevitably many mean free paths thick. This makes it an excellent neutron absorber. The bomb casing is not compressed in the same sense, and would have to be extremely thick and heavy to capture many neutrons.

4.4.4 Implosion Systems

The discussion above (and in Section 3.7.5 Methods for Extreme Compression) it has been made clear that efficient compression requires creating and maintaining a relatively low pressure for a relatively long time, with the pressure rise accelerating rapidly near the end of the compression process. But so far, I have not discussed at any length how this might be accomplished.

To make this problem clear consider the required duration of relatively low pressure. For a typical fuel layer thickness of 2 to 8 cm (depending on weapon size), it would take the weak initial shock (travelling at ca. 100 km/sec) something like 200-800 nanoseconds to traverse it. During most of this time, perhaps 80% of it, the pressure at the fuel surface can be permitted to rise no higher than a few tens of megabars. The remaining pressure increase - to a value perhaps a thousand times higher than the average pressure of the initial shock - can occur no sooner than this final 20%.

But the source of this pressure - the primary - typically generates its energy output on a much shorter time scale. This time scale is determined by the length of the multiplication interval, 1/alpha, which may be no more than a few nanoseconds. Within a time period of a few times 1/alpha, say 3-4 multiplication periods, >98% of the fission reactions occur and we can think of essentially all of the fission explosion occurring during this time. Thus nearly all of the energy and excess neutrons produced by the primary are released within perhaps 10-15 nanoseconds for a pure fission primary, and as little as 3-4 nanoseconds for a fusion boosted primary.

Clearly some cleverness is required to stretch out the rate at which this brief burst of energy arrives at the fusion fuel.

4.4.4.1 Techniques for Controlled Implosion

A number of techniques for doing this can be identified which may be used alone or (probably more typically) in combination to achieve the desired pressure vs. time history.

4.4.4.1.1 Release Waves

The development of a release wave when the ablation shock completes its passage through the tamper (see Section 4.4.3.2.1 The Ablation Shock above) is an inherent feature of radiation implosion which significantly contributes to achieving efficient compression. The release wave converts the sudden intense pressure jump of the ablation shock front into a lower pressure, higher velocity shock in the fusion fuel which is followed by a gradient of increasing pressure.

We previously estimated the ablation shock velocity for Mike at 160 km/sec, and 570 km/sec for the W-80. The release wave driven shock must be even faster. This indicates the release wave driven shock will be much faster, and its pressure much higher, than the relatively weak 50 megabar bar shock (travelling at 100-200 km sec) described earlier. We can conclude then that unless the ablation shock pressure is very low, this mechanism does not by itself reduce the shock jump sufficiently to give efficient compression.

4.4.4.1.2 Standoff Gaps

A standoff gap is a void between the fusion fuel and the tamper. The effect of a standoff is to allow the release wave to unload to zero pressure and full escape velocity (see Section 3.6.1.1 Release Waves), converting the internal energy of the gas entirely into kinetic energy. The forward edge of the wave then runs far ahead of the bulk of the imploding tamper without heating any fuel in the process.

When it reaches the fusion fuel, the release wave will be decelerated and begin piling up at the void/fuel interface, driving a low pressure shock into the fuel. As the rest of the release wave arrives the pressure keeps climbing, driving a compression wave of increasing strength.

The velocity of the release wave front is given by:

u_escape = (2*c_s)/(gamma - 1) + u_particle

where c_s is the speed of sound in the shocked tamper, and u_particle is the velocity acquired by the tamper from the shock wave. Thus for a perfect monatomic gas, with a gamma of 5/3, this velocity is equal to three times the speed of sound in the shock compressed gas plus the velocity the gas acquired by passing through the shock front. In a dense tamper under extreme pressure the effective value of gamma may be significantly lower than 5/3 due to ionization effects, making the escape velocity four or more times faster than the speed of sound.

If the tamper were not accelerating, then the larger the standoff gap the greater the elapsed time between the arrival of the release wave and bulk of the tamper, which is desirable for efficient compression. But the tamper is actually accelerating, so in time it will tend to catch up with the release wave front. For a given geometry, ablation pressure, and tamper mass, there is an optimum standoff that will maximize the elapsed time.

The use of use of a standoff seems to have been the major (perhaps only) method for creating the desired compression wave in Mike, the first radiation implosion device ever tested. From the available specifications, we can estimate that the standoff may have been in the order of 25 cm, with a fuel mass radius of 20 cm. Calculating u_escape at around 600 km/sec (gamma = 1.5,

c_s = 100 km/sec), the elapsed time between the arrival of the release wave and the rest of tamper at the initial fuel surface radius would be about 300 nanoseconds. At an average shock velocity of 200 km/sec, the initial shock could traverse 8 cm of fuel before the tamper finally catches up, far enough to efficiently compress 64% of the fuel.

4.4.4.1.3 Compartmented Radiation Cases

A second technique is to divide the interior of the weapon into two compartments that separate the primary and secondary. A barrier between the compartments made of opaque (high-Z) material controls the rate at which energy flows from the primary to the secondary.

Since a small amount of energy is needed to begin the implosion, the barrier would have tiny apertures (narrow slits perhaps) that would allow photons to enter the secondary compartment at a slow rate. The barrier material ablates away, driving an ablation shock through the wall. The ablation shock is luminous (though much less intense than the unobstructed flux from the primary compartment would be) so when it arrives at the opposite side, a significant additional thermal flux into the secondary compartment would occur.

By far the largest increase in radiation flux would occur when the ablation front arrived at the opposite side of the barrier (i.e. when it completely ablates away). Then radiation at the full temperature of the primary compartment would flow into the secondary compartment.

Of course the barrier would be driven forward at a very high velocity by the ablation shock, and preventing it from damaging the secondary would be a significant problem. One possible technique for addressing this problem would be to place a shield made of X-ray transparent low-Z material (lithium, beryllium, or boron for example) between the barrier and the secondary to absorb the impact of the barrier remnants.

Many variations on this idea are possible. Varying the thickness or the composition of different parts of the barrier could provide a more carefully tailored release of energy. Thermal energy could be diverted into "radiation bottles" by unimpeded flow through a duct or pipe before release to the secondary. Multiple barriers or baffles could be used to control the rate of energy flow.

4.4.4.1.4 Modulated Primary Energy Production

The idea here is to tailor the energy production in the primary so that the desired pressure-time curve is produced directly. The functional form of fission energy release (an exponential function) actually does match the desired functional form of the pressure-time curve fairly well. The problem is that the time constant of a reasonably efficient fission system is simply to short. By the time a low pressure shock created by an early stage of fission has propagated a substantial distance (a few millimeters, say) the intense shock from the final stages of fission will have caught up with it. If the value of alpha is reduced to the point where the rate of increase is tolerably slow (10-20 per microsecond), the core has time to disassemble without producing much energy - leading to a very inefficient primary.

It may be possible to use fusion boosting to overcome this problem. Since boosting can be initiated at a fairly low fission yield and accelerates as the temperature rises, it may be possible to use boosting to still achieve high efficiency. Boosting would kick in after the slow, low pressure phase and drive the rapidly rising high pressure end of the curve.

A design of his kind would have several advantages. The low alpha of the fission process would mean that a small quantity of fissile material and/or weak compression would be adequate for the primary, leading to light and compact primaries. The requirements for radiation containment would be reduced as well, leading to reduced overall weapon weight.

A disadvantage is that the idea could not be extended to weapons of unlimited yield. Larger yields require thicker fuel layers, slower initial compression, lower alpha values in the primary, and reduced fission yield. The approach would fail (if it can be made to work at all) when the reduced alpha value allows the primary to disassemble before initiating the boosting process.

A primary using this approach would be designed with a beryllium reflector, but with no tamper between the fissile material and the reflector so that radiation escaped as readily as possible. U.S. primaries are known to contain plutonium bonded directly to beryllium, suggesting this design approach.

4.4.4.1.5 Multiple Staging

In weapons with more than two stages, the efficient compression of tertiary (or, in principal, later stages) can be conveniently arranged with the aid of the sequenced energy release of the earlier stages. This is fundamentally the same general idea of modulated energy release just described, using a different mechanism.

The secondary stage releases much more energy than the primary (as much as 200 times more has been demonstrated, but more typically 10-50 times more), and does so hundreds of nanoseconds later.

A portion of the primary energy can be used to create an initial low pressure shock in the tertiary stage, even as it compresses the secondary. The third stage which would generally be larger and have a greater radius that must be traversed by the initial shock, requiring a longer compression period in any case. The sudden burst of energy from the secondary would be quite effective in creating the rapid rise in pressure desired at the end of the tertiary compression period.

This technique is obviously of limited general usefulness, since only relatively large weapons would have three stages (all known three stage tests have been in the megaton range, very few three stage designs appear to have been actually fielded).

4.4.4.1.6 Selection of Pusher Materials

Another possible technique for creating a time varying pressure in the fuel is to modify the ablation process itself. The amount of ablation pressure generated by radiant heating depend on the properties of the material being ablated.

If the ablation surface has a very high atomic number, then the ablated gas will still be quite opaque to X-rays. This means that the radiation will have to reach the ablation front by diffusion - each X-ray being captured and re-emitted multiple times. Radiation diffusion is a relatively slow process. Also, the hot ablated gas will radiate energy back into the radiation channel, reducing the net flux reaching the ablation front.

A lower Z material, which completely ionizes at the radiation channel temperature, will become nearly transparent to X-rays when heated. The X-ray flux will thus reach the cold ablation surface unimpeded. Neither the cold surface, nor the hot gas, will radiate significant amounts of energy back into the channel so the thermal energy will be absorbed by the ablator very rapidly (with a correspondingly high mass loss rate).

The effective particle mass of a completely ionized low-Z material will be much lower than that of a partially ionized high-Z material. This gives a higher escape velocity, and a larger ablation pressure per unit of mass lost.

These factors give the designer a range of materials and effects to choose from to tailor the ablation rate and pressure. Using multiple layers of different materials offers the possibility of creating a time-varying ablation pressure even with constant radiation temperature.

No information is available indicating that this technique has been ever actually been used.

4.4.4.2 Radiation Containment and Transport

The thermal radiation that drives the implosion process must be kept from escaping until it has completed its work. This is the function of the radiation case.

The radiation must also be transported rapidly and effectively to the secondary (when it is time). The conduit through which the radiation flows is the radiation channel.

4.4.4.2.1 Radiation Case

The radiation case may be integral with, or identical to, the external bomb casing; or it may be a separate component nested inside of the bomb casing. There may in fact be more than one radiation case, especially in a multi- compartment design. Radiation cases may also be nested inside each other, to provide different degrees of confinement during different phases of bomb operation.

To fulfill its role, the wall of the radiation case must be highly opaque to the radiation that fills it to minimize the rate at which energy is lost to the wall (see Section 4.4.3.2 Opacity of Materials in Thermonuclear Design). In general the radiation case will either be a lining of the external bomb casing, or will be entirely separate from it.

It is inevitable that the wall of the case will ablate away, just as the secondary pusher does, and will generate a high pressure shock that blows the wall outward. To minimize the rate of case expansion, the casing wall should also have a very high mass density. The DOE has reported that materials with atomic numbers higher than 71 have been used as radiation case linings. Uranium, lead and lead-bismuth alloy are known to have been used to line radiation cases. Tungsten, or tungsten-rhenium alloys (such as the thin plasma deposited tungsten-rhenium coatings developed at the Kansas City Plant) are also good candidates for this purpose. Mercury, thallium, and gold have been used in thermonuclear weapons - possibly for this purpose.

The radiation temperature around the secondary needs to be maintained until the secondary collapse is complete or nearly so, otherwise the outer wall of the imploding tamper will decompress and begin expanding, reducing its ability to confinement the thermonuclear reaction. This defines the length of time that the radiation case must maintain its integrity (at least in the part of the weapon where the secondary is located).

The rate of energy loss into the wall probably remains more or less constant until the ablation shock arrives at the outer surface of he case. Once this occurs, the wave of pressure release will travel backwards to the inner wall relatively rapidly. When this release wave reaches the ablation front, the rate of energy loss will rapidly increase - ending the useful life of the casing. The arrival of the Marshak wave front of the ablation shock at the outer surface of the casing was an important diagnostic in early thermonuclear weapon tests.

In two compartment weapon designs a separate casing is placed around the primary. Since the pressure here is initially much higher (and persists longer) than around the secondary, a thick walled sphere of uranium is used to provide an especially opaque and dense case.

For weapons that use the soft X-ray kill mechanism (e.g. high altitude ABM or space-based interceptors), a radiation case that is transparent to the more energetic X-rays produced by the secondary is desirable. Since the average photon energy during implosion is only 2 KeV or so, and the bulk of the energy emerging from the secondary is carried by photons with energies >>20 KeV, this should not be too hard to arrange. In fact with a lining of sufficiently low Z, the hot photon flux should be capable of completely stripping the nuclei of electrons through photo-ionization, rendering it essentially transparent ("bleaching it").

4.4.4.2.2 Radiation Channel

The outer wall of the radiation channel is the radiation case. The inner wall is generally the pusher of the secondary. The thermal energy released by the primary is conducted down the channel by diffusion - a given region of the wall is heated by radiation emitted by hotter regions of the wall closer to the primary, and in turn re-emits radiation to heat regions of the wall that are farther away.

The rate of energy flow any point in the channel can be modelled by the diffusion equation:

J = -((photon_MFP * c)/3) * (energy_density_gradient)

where J is the energy flux (flow rate), photon_MFP is the mean distance a photon travels between emission and capture, c is the speed of light, and energy_density_gradient is the rate at which energy density changes with distance along the channel. The temperature also changes along the channel with energy density, but since temperature is proportional to the fourth root of energy density, the gradient here is much smaller.

If the radiation channel is transparent, then the photon mean free path is the average distance down the channel a photon will travel between emission and absorption. This is determined by channel geometry (plane, cylindrical, spherical, etc.), and the width of the channel. Transport is faster along a straight channel than a curved one, and faster along a wide channel than a narrow one. For an arbitrary small patch of channel wall, the proportion of energy emitted at an angle theta from the normal vector is Cos(theta). The distance the energy will travel down the channel before absorption is **Tan(theta)*channel_width**. The average of **Cos(theta)*Tan(theta)*channel_width** is simply **channel_width**. As long as the mean free path in the material filling the channel is substantially less than **channel_width/Cos(45 degrees)**, about 1.4 times the channel width, this will be the effective mean free path down the channel.

Since c, the speed of light, is very large the rate of transport tends to be extremely fast. The energy density will thus very rapidly come into equilibrium, as long as the maximum distance between two points in the channel, as measured in photon mean free paths, is not also a very large number. Even when energy is flowing into the channel, the energy density gradient will remain quite small. If energy is not flowing into the channel, any irregularities will rapidly disappear.

The ablation of the channel walls interferes with the need to maintain a transparent channel. The high-Z material lining the channel produces a high velocity gas as it escapes from the channel wall, and it accelerates further as it expands into the channel. Even at relatively low densities this gas is quite opaque, and it has the effect of rapidly collapsing the radiation channel until it is completely blocked.

This process can be combated by filling the channel with a transparent gas to hold back the ablating walls. It is impossible to hold back the ablating material completely, but the highest velocity ablation exhaust is at low pressure and is relatively easy to contain. As the gas-filled channel closes, its pressure increases as well making it more resistant to further collapse.

Radiation channels are typically filled with a dense plastic foam such as polystyrene, that has been "blown" (foamed) with a hydrocarbon gas (pentane for example). The channel is thus filled only with carbon and hydrogen. The ionization energy of the last electron in carbon is 490 eV (hydrogen's ionization energy is a mere 13.61 eV), which corresponds to the average particle energy at a temperature of 5.7 x 10^6 K. As the radiation channel approaches this temperature the foam will become completely ionized and nearly transparent to thermal radiation. Polyethylene wall linings have been used instead of plastic foam (in Mike for example) although unless the casing is flushed with a low-Z gas, the higher ionization energies of nitrogen and oxygen may cause significant absorption.

Note: that this foam *does not* generate the pressure that causes implosion.

4.4.4.3 Avoiding Fuel Preheating

Compressing fuel efficiently to high densities requires that the fuel have relatively low entropy. At the start of the compression process, a relatively small amount of heat will increase the entropy significantly and reduce the efficiency of the entire compression process, which is why the initial shock pressure must be carefully controlled. Fuel preheating can also occur from the radiation emitted by the primary. A high-Z radiation shield is used to prevent the X-ray flux from directly heating the fuel in cylindrical designs, but the neutron flux from the primary can also cause significant preheating. This can not only reduce the achievable compression, but due to uneven heating in the fuel it could disrupt the symmetry of implosion as well (a potentially even more serious problem).

If we assume that 1.5 neutrons escapes from the primary for each fission occurring there, then up to 3 MeV (about 1.5% of the total yield since the average fission neutron has 2 MeV) could potentially be carried away from the primary by neutron kinetic energy. The portion of this that would be deposited in the fuel depends on the area of the fuel presented to the flux, the distance from the primary, and the effects of materials between the core and fuel in absorbing neutron energy. It is difficult to see how any more than 5% or so of the flux could be intercepted by the fuel, and is likely to be much less than this. The average energy can be expected to be significantly less also, due to moderation by the beryllium reflector and high explosive.

In a weapon that has a fuel mass/primary yield ratio of one kg per kt, intercepting 1% of the neutron kinetic energy emitted by the primary core is still some 6x10^12 erg/g. The problem of preheating is especially serious in lithium-6 containing fuel, since the Li-6 + n reaction releases 4.8 MeV in addition to whatever kinetic energy the incident neutron possesses. Spherical secondaries are more likely to be prone to this problem than cylindrical ones, since they present a larger surface area. The Morgenstern device (designed by Edward Teller) that fizzled in the Castle Koon test reportedly had a spherical secondary, and failed due to neutron preheating effects.

Neutron preheating can be avoided by attenuating the neutron flux with boron-10, the best available fast neutron absorber. Boron carbide (B6C) is known to have been incorporated into thermonuclear weapons, possibly for this reason. Such a neutron thermal shield would be incapable of stopping the neutron flux from reaching the secondary completely, but at most a one or two order of magnitude reduction should be sufficient to render preheating insignificant. With Z equal to 5, boron is highly transparent to thermal X-rays and would not interfere with radiation transport.

4.4.5 Fusion Stage Nuclear Physics and Design

4.4.5.1 Fusionable Isotopes

The important thermonuclear reactions for weapons are given below:

1. D + T -> He-4 + n + 17.588 MeV (n kinetic energy is 14.070 MeV)

2. D + D -> He-3 + n + 3.2689 MeV (n kinetic energy is 2.4497 MeV)

3. D + D -> T + p + 4.0327 MeV

4. He-3 + D -> He-4 + p + 18.353 MeV

The first fuel ever considered for a thermonuclear weapon was pure deuterium (reactions 2 and 3, which are equally likely). This is primarily because deuterium is a relatively easy fuel to burn (compared to most other candidates), is comparatively abundant in nature, and is cheap to produce. In fact, no other fuel has this same combination of desirable properties.

Only one other fusion fuel is easier to ignite - a mixture of deuterium and tritium (reaction 1). At moderate thermonuclear temperatures, the T-D reaction is 100 times faster than D-D combustion. Unfortunately, tritium does not occur in nature in useful amounts, and is very costly to manufacture.

The cheapest method of making it industrially is to breed it in reactors, where it competes with plutonium production. For each neutron absorbed in the reactor for isotope production, either one atom of tritium or one atom of plutonium can be produced. Since fusing an atom of tritium produces 17.6 MeV of energy, compared to 180 MeV from fissioning an atom of plutonium, the cost of the energy represented by tritium is ten times that of plutonium. Worse still, it decays at a rate of 5.5% annually so simply maintaining an inventory of tritium is expensive. Unless the energy output of tritium can be magnified greatly, or its effective cost greatly reduced, it is uneconomical to use it weapons.

Tritium can be produced in situ from other reactions in a weapon. Deuterium - deuterium combustion, for example, produces tritium naturally through reaction 3. In fact far more energy is produced in D-D fusion from fusion of the tritium byproduct than from the D-D reaction itself. Since the D-T reaction rate is far higher than the D-D rate, and there is always a large excess of deuterium, nearly all the tritium produced is burned up.

The helium-3 + D reaction (reaction 4) is even more energetic than the D+T reaction, but it is harder to ignite. The cross section is much lower than the D+D cross section at temperatures below 200 million degrees K. Helium-3 is not found in useful amounts on Earth but, like tritium, it is produced as a by-product of D-D fusion. Reaction 4 only becomes important with pure deuterium fuel when a significant amount of deuterium has burned up (about 25%). At this point, the temperature has risen to about 250 million degrees K, where the cross section for reaction 4 begins to exceed that of 2 and 3 combined. Also, at this point the concentration of He-3 has built up to be a significant proportion of the fuel mass. The conversion of He-3 to tritium through neutron capture competes with the build-up of helium-3 however (see reaction 10 below).

The net effect of reactions 1-4 together is:

6 D -> 2 He-4 + 2p + 2n + 43.243 MeV

Of the two neutrons produced, one is high energy (14.07 MeV) and one is moderate energy (2.450). The ratio of high energy neutrons produced to deuterons consumed (or energy produced) is significant for driving fast fission reactions. If the He-3 is converted to tritium instead of being burned directly, the net reaction is the same with the exception that two high energy neutrons are produced.

There are other fusion reactions that occur between these isotopes (T+T and He-3+T for example), their reaction products, or with other materials commonly mixed with the fusion fuel (like lithium isotopes), but the reaction rates are too low to be significant.

4.4.5.2 Neutronic Reactions

The neutrons released by reaction 1 and 2 can be put to use in several ways. They can be allowed to escape the weapon to serve as one of the destructive weapon effects. They can be used to cause fission (perhaps in cheap non- fissile material like U-238 or Th-232), thus releasing additional energy. Or they can be used to manufacture more fusion fuel to enhance the fusion reaction. Both of these last two possibilities are commonly incorporated into modern weapons. The use of neutrons as a distinct weapon effect is usually only important in special designs (neutron bombs).

The two reactions that have been used to manufacture fusion fuel are:

5. Li-6 + n -> T + He-4 + 4.7829 MeV

6. Li-7 + n -> T + He-4 + n - 2.4670 MeV

Both produce tritium which burns rapidly, producing more neutrons.

Lithium-6 is a relatively uncommon isotope in nature (7.42% of natural lithium) and must be enriched before reaction 5 can be used to best effect. The Li-6 + n reaction has a significant cross section at all neutron energies, but it has a large cross section below 1 MeV with a peak of 3.2 barns at 0.24 MeV. At higher energies endothermic spallation reactions tend to occur instead, above 4 MeV the neutron is far more likely to split the Li-6 nucleus into He-4 and D.

Lithium-7 constitutes the bulk of natural lithium (92.58%). The endothermic Li-7 + n reaction does not occur at all if the neutron energy is less than the energy deficit, and is only significant with neutron energies above 4 MeV. Above 4.5 MeV Li-7 has a much larger cross section from breeding tritium than does Li-6.

There are a number of side reactions that can also occur in fusion fuel, especially with the very energetic 14.07 MeV fusion neutrons, which can cause spallation or fragmentation of target nuclei due to their enormous kinetic energy. Among these are:

7. D + n -> p + 2n - 2.224 MeV

8. Li-6 + n -> He-4 + D + n - 1.474 MeV

9. Li-6 + n -> He-4 + p + 2n - 3.698 MeV

10. He-3 + n -> T + p + 0.7638 MeV

11. Li-7 + n -> Li-6 + 2n + -7.250 MeV

In deuterium and Li-6D fuel reactions 7 and 9 are significant in causing a modest amount of neutron multiplication (10-15% amplification of 14 MeV neutrons), and aiding in the rapid attenuation of highly energetic neutrons. Reaction 10 is especially important in pure deuterium fuel where helium-3 produced by the D-D reaction is the only significant neutron absorber.

With respect to the total energy release, and the composition of the final products, the net effect of reactions 10 and 1 together is exactly the same effect as reaction 4. That is, converting He-3 to T, then fusing it with D is the same as fusing He-3 with D directly. The energy is not distributed over the reaction products in exactly the same way however. Reaction 10 consumes a neutron but this may be a very low energy neutron (in fact it most likely will be given the very large cross section below 0.5 MeV - up to 5 barns). Reaction 1 produces a 14.07 MeV neutron. In effect a very high energy neuron is exchanged for a low energy one. This can change the ratio of high energy neutrons to deuterons consumed from 1:6 (implied by reaction 1-4) to as low as 1:3, greatly augmenting fast fission.

Note also that when the above lithium-7 reaction (reaction 11) is combined with reaction 5, the net effect is exactly the same as reaction 6.

Since fusion fuel contains a very high density of very light atoms (like deuterium) with good scattering cross sections, we should expect neutrons entering the fuel to be rapidly moderated.

In pure deuterium fuel moderation takes only 9 collisions to fully thermalize 14.07 MeV neutrons ("thermal" here means on the order of 20 KeV), a process essentially complete in 0.25 nanoseconds at a fuel density of 75 g/cm^3. Deuterium's absorption cross section for neutrons at 20 KeV is only about 1 microbarn, given the neutron velocity (2x10^8 cm/sec) and atom density (2.25x10^25 atom/cm^3 at 75 g/cm^3) the lifetime of a thermalized neutron is about 220 microseconds (for this reaction). The cross section for reaction 10 at 20 KeV is 5 barns so when 10% of the fuel has burned (creating an He-3 concentration of 2%), the lifetime of a thermalized neutron before He-3 capture will have dropped to 2.2 nanoseconds. Thus the formation of tritium through neutron capture by deuterium cannot play any significant role, as some have suggested.

Rapid moderation occurs in lithium deuteride fuel as well. Reaction 5, the production of tritium from lithium-6 has a very large cross section peak at 246 KeV (8.15 barns). It averages only 0.77 barn from 0.02 to 0.15 KeV and 1.1 barns from 1 to 14.1 MeV. Multi-group neutron calculations show that in Li-6D fuel at a density of 200 g/cm^3 about half of all tritium production occurs with neutrons moderated to the range of 0.15-1.0 MeV. 50% of 14.07 MeV neutrons are absorbed to form tritium within 0.075 nanoseconds after emission, rising to 70% at 0.15 nsec. Most of the rest become thermalized with a lifetime of 0.40 nsec before capture.

The above observations for 14.07 MeV neutrons remain valid for the 2.45 MeV D-D reaction neutrons, except that fewer collisions are required for the moderation. Clearly fusion neutrons give their energy up very quickly to the fusion fuel, and relatively few escape the fuel without undergoing substantial moderation. We can also conclude that the production of tritium from lithium-6 is a rapid, efficient process.

4.4.5.3 Fusion Fuels

These are fuels that produce energy primarily through charged particle reactions, driven by thermal kinetic energy. Neutron reactions often play important anciliary roles.

4.4.5.3.1 Pure Deuterium

Deuterium is an inexpensive fusion fuel, consisting on the order of $100/kg to manufacture, with an effectively unlimited supply. Its major disadvantage is that it is a gas at normal conditions, requiring extreme cold to liquefy it (to below 23.57 K). It has the additional disadvantage that it is a low density liquid - 0.169 (or 0.0845 moles/cm^3). This low density, combined with the necessity of extremely efficient insulation implies a large volume for any weapon using pure deuterium as a fuel (to say nothing of the cost, weight, and complexity of the cryogenic equipment needed for handling it, and keeping it cold).

Deuterium has a high energy content however, 82.2 kt/kg with complete thermonuclear combustion. It also produces a large excess of neutrons per unit of energy released, one neutron for each 21.62 MeV of reaction energy. The net reaction is:

6 D -> 2 He-4 + 2 p + 2 n + 43.24 MeV

Pure deuterium has been used in at least one thermonuclear test - Ivy Mike, the first radiation implosion design ever tested. The fact of this test conveniently demonstrates that thermonuclear energy release in weapons does not require tritium breeding neutronic reactions, but can be driven by the D+D reactions alone.

4.4.5.3.2 "Dry" Fuels (Lithium Hydrides)

It would be more convenient if deuterium could be incorporated into weapons in the form of a stable chemical compound with more convenient physical properties than the low boiling point elemental form. A suitable compound would be the hydride of a light element, which would give a fairly high deuterium content by weight.

While there are several compounds that fit this description, it was realized quite early in both the US and Soviet Union that one compound in particular was uniquely suited for this role - lithium deuteride. Even more important than its high deuterium content (22.4-25% by weight), and high atom density (0.103 moles D/cm^3, higher than in liquid deuterium!), is the fact that lithium isotopes can also provide additional fusion fuel. By capturing neutrons generated as fusion byproducts, reactions 5 and 6 produce highly combustible and energetic tritium. Reaction 5 also

produces significant amounts of energy directly from neutron capture. Probably all fusion devices since Mike have used lithium hydrides of varying isotopic composition as fusion fuel.

4.4.5.3.2.1 Enriched Lithium Deuteride

The most desirable fuel is pure lithium-6 deuteride since it has the highest energy content per kilogram: 64.0 kt/kg. The net reaction is a combination of reactions 1 and 5:

Li-6 + D -> 2 He-4 + 22.371 MeV

There are a few considerations that must be addressed before this reaction will work. First, the neutrons produced by reaction 1 are too energetic to direct drive reaction 5 efficiently - they must undergo a few collisions to moderate their energy. Also, there must be an initial source of neutrons or tritium to drive reaction 5 before reaction 1 can occur. The overall cycle does not breed neutrons.

Some open literature sources assert that reaction 5 is driven by neutrons produced by fission reactions in the trigger, the spark plug, or the tamper.

The first of these suggested sources can be easily disposed of as a possibility. If neutrons from the primary were to breed a significant amount of tritium, severe neutron preheating problems would result.

A number of arguments can be offered against the other possibilities. The most obvious is that the net Li-6 + D reaction does not produce spare neutrons (although a small excess of 10-15% might be produced though n->2n reactions with the fast neutrons). Since only a relatively small proportion of the neutron excess can actually cause fast fission in U-238 (due to moderation, inelastic scattering, and absorption), for a lithium deuteride fueled bomb to produce substantial energy through fast fission some other type of fusion reaction must provide the neutron excess.

It can also be observed that it is very difficult to construct a scheme that will permit neutrons from the spark plug to play a major role. "Clean" weapon tests have been conducted that obtained as much as 98% of their yield from fusion reactions. Even if all the fission were due to the spark plug (allowing us to neglect the trigger), some 9000 MeV of fusion would have to result from the neutrons released by each 180 MeV fission (producing fewer than 2 excess neutrons). This implies a process of neutron recycling (a neutron is absorbed to form tritium, tritium fuses to release a neutron) some 200 reactions long. It is likely that neutron leakage would quench this chain long before it got this far.

We can conclude that the source of neutrons to prime the pump for the Li-6 + D reaction, and provide the neutron excess for fast fission, is D-D fusion. This is natural since D-D is capable of supporting energetic fusion power production by itself. Even without tritium breeding, D-D fusion burns to effective completion in approximately 20 nanoseconds anyway, producing a staggering number of neutrons in the process. The presence of lithium-6 means that these neutrons are soaked up, producing even faster burning tritium and accelerating the combustion process.

To summarize: D-D initially dominates the fusion process, but as the neutron concentration builds up it is quickly superseded by Li-6 bred tritium fusion as the prime power producer.

Since lithium-6 constitutes only 7.42% of natural lithium, making use of lithium-6 deuteride requires enriched lithium. Varying enrichments can be used, but the higher the better. The U.S. has produced and used enriched lithium containing 95.5% Li-6, 60% Li-6, and 40% Li-6. It appears that 95.5% has made up the bulk of U.S. production (442.2 tonnes total).

4.4.5.3.2.2 Natural Lithium Deuteride

Lithium-7 can serve as a fusion fuel also, either in partially enriched lithium or in natural lithium. The unexpected contribution from Li-7 in the Shrimp device tested in Castle Bravo (which used 40% Li-6, 60% Li-7 fuel) caused it to exceed predictions by a factor of 250% (to 15 Mt, the largest U.S. test ever). Natural lithium has also been used successfully as a fusion fuel in tests (it was used in the 11 Mt Castle Romeo test for example), and fielded weapons.

The contribution from lithium-7 is primarily due to reaction 6, which is quite significant above 4.5 MeV. In natural lithium the probability that an emitted 14.1 MeV neutron will breed tritium by reacting with Li-7 capture is around 50% (failure to appreciate this fact led to the Castle Bravo disaster). Since reaction 6 does not actually consume a neutron (a low energy neutron is one of the reaction products), in effect the neutron is available to also react with any lithium-6 that is present to breed additional tritium.

The net reaction for lithium 7 is a combination of reactions 1 and 6:

Li-7 + D -> 2 He-4 + n + 15.121 MeV

for an energy content per kilogram of 38.5 kt/kg. Pure Li-7 is unlikely to be used of course. The light isotope in natural lithium will react as well yielding 40.3 kt/kg.

The fact that Li-7 actually does breed additional neutrons may be of significance in enhancing energy production through tamper fission.

4.4.5.3.3 Speculative Fuels

Deuterium can be combined with any of the other light elements (except helium) to form chemical compounds that could probably be used successfully as fusion fuels. Most of these have higher deuterium contents than lithium hydride, and all of them are easier to store than liquid hydrogen. Below is a list of representative compounds that can be formed with each of the elements from Z=4 (beryllium) to Z=8 (oxygen).

Given the convenience of lithium deuteride's physical properties, the cheapness of lithium, and the high energy content resulting from the nuclear reactions lithium undergoes, there is no reason to prefer any of the fuels below to LiD. It is quite unlikely that any of them have ever been used.

Deutero-ammonia (ND3) was seriously considered as a fusion fuel for the first radiation implosion device (Mike) in preference to liquid deuterium and natural lithium deuteride.

I have not investigated the possible nuclear reactions that the compounding agents below might undergo but one of them, boron-10, does have an exothermic reaction with neutrons which would add modestly to energy production, which produces a usable fusion fuel besides: B-10 + n -> Li-7 + He-4 + 2.79 MeV

The possibility of using ordinary heavy water as a fusion fuel is interesting given its comparative availability and good physical properties.

Formula	Deut. Conc.	Properties
B2H6	35.6%	Bp -92 C, cryogenic cooling required
BeD2	30.8%	solid, stable to 125 C
BeD2.B2H6	33.2%	solid
CD4	40.0%	Bp -164 C, cryogenic cooling req'd (same as LNG - liquid natural gas)
C2D6	33.3%	Bp -88 C, cryogenic cooling or high pressure required
C3D8	30.8%	Bp -42 C, storable at room temp under modest pressure
D2O	20.0%	Bp 101.5 C, liquid, no special handling
ND3	30.0%	Bp -30.9 C, storable at room temp under modest pressure

4.4.5.4 Fusion Tampers

The jacket surrounding the fusion fuel of the secondary is often called upon to perform quadruple duty:,

1. It provides reaction mass as an ablator to drive the radiation implosion;

2. It acts an inertial mass to confine the fusion fuel during the burn;

3. It acts as a radiation container to prevent loss of heat during the burn;

and

4. It also acts as an energy producing fuel by reacting with the neutrons produced by the fusion reactions.

The first three functions are essential for successfully releasing energy from the secondary. The fourth function is optional (fission energy release that is), and it is convenient for weapon designers that materials that meet this last requirement also satisfy the first three quite well.

A variety of materials are available that can meet these requirements. Functions 1 and 3 basically call for a material with a high opacity at high temperatures and pressures - a high Z element. Function 2 ideally calls for an material that will be compressed to very high density, which is provided by a material that only partially dissociates under high pressure - also a high Z element. Function 4 calls for a fissionable isotope (at least with fast neutrons) - the only plausible candidates for this (uranium and thorium) are the two highest Z natural elements available in substantial quantity.

Although the use of a fissionable tamper is optional, fission of the tamper provides the majority of the energy released in most thermonuclear weapon designs. This is because the energy produced is essentially "free". The tamper mass is required in any case, so extracting energy from it increases the bomb's yield for the same weight. The cost of the fissionable tamper material, in terms of its available energy content, is also generally much lower than the other materials used in the bomb, so the cost increase of using a fissionable tamper over a cheaper non-fissionable tamper is small.

A high fission yield produces a large amount of radioactive fallout. Weapons using fissionable tampers are thus considered "dirty". If this is considered undesirable in the weapon, then a non-fissionable tamper may be used. The penalty of course is a lower yield to weight ratio, and a more expensive arsenal to deliver the same amount of destructive power.

On the other hand, an increase in fallout relative to yield may be desired, perhaps with a custom tailored range of half-lifes. Using a tamper that produces highly radioactive byproducts when irradiated by neutrons holds this possibility.

The required mass of the pusher/tamper is set by the several requirements. The need to provide sufficient ablation mass (much of the pusher/tamper mass is lost during implosion through ablation), sufficient inertia to provide the momentum need for fast fuel implosion, sufficient inertia to confine the fuel during burnup, and sufficient thickness to capture a high percentage neutrons produced (if this is desired).

If we assume maximally efficient implosion, then 75% of the pusher/tamper mass will be lost through ablation. To provide substantial inertial for driving implosion and confinement, we would also like the remaining tamper mass to significantly exceed the fuel mass. This puts the probable ratio between the pusher/tamper and fuel mass in the range of 8-16:1 (the Mike device had a ratio of something like 80:1).

The objective of capturing most of the neutrons should not difficult to achieve either. At densities of around 500 g/cm^3 for uranium , the thickness of tamper required to reduce the flux to 1/e (36.8%) of the initial value is no more than 0.5 cm.

The high proportion of the pusher/tamper that should ablate for efficient implosion, combined with the difficulties in estimating or measuring the high temperature LTE opacities, created a significant problem for early designers. If too much of the pusher/tamper ablated, complete failure of the secondary could result through insufficient confinement or even burn-through before implosion was complete. If not enough ablated, the implosion might be too inefficient to give a good yield. This necessitated conservative design, and may be the explanation for the apparent failures of the first British tests (the Green Granite and Purple Granite devices).

4.4.5.4.1 Fissionable Tampers

The advantage in using a fissionable material as the fusion tamper was recognized very early, many years before the first thermonuclear test. The highly energetic neutrons produced by fusion are capable of fissioning isotopes that are normally considered non-fissile, like U-238 and Th-232.

The early high yield fission weapon designs all used natural or depleted uranium as the tamper material. At the time large, inexpensive, high yield weapons were the main design objective so a cheap fissionable material was necessary. U-238 can only be fissioned by neutrons above about 1.5 MeV however. A 14.1 MeV neutron can undergo up to three average collisions with deuterium and still have sufficient energy. the 2.45 MeV D-D fusion neutrons, up to half the fusion neutrons produced in Mike, cannot be scattered even once and still be able to fission the tamper. Most of the fusion neutrons produced were thus unable to fission the tamper after they escaped the fusion fuel mass.

If all of the excess neutrons had caused fission, then the expected fission fraction for Mike (assuming no subcritical multiplication took place in the tamper) would be 89.3% (10.4-1 fission-fusion energy production ratio), instead of the observed 77% (3.35-1 ratio). Isotopic analysis of Mike fallout shows a very high percentage of the tamper material that was not fissioned was transmuted to higher isotopes of uranium by these slower neutrons (~93% of the U-238 in the inner most layers was transmuted).

If the uranium tamper is significantly enriched in U-235 however, a much higher percentage of the neutrons released can be harnessed since this isotope is fissioned by neutrons of all energies. The superior fissile properties of U-235 boost yield in other ways also. It has a higher fission cross section than U-238 even for fast neutrons, so a thinner tamper can be equally effective in capturing these as well. U-235 also can achieve much higher subcritical multiplication factors in the tamper - making use of the fission neutrons to cause more fissions.

Using enriched uranium as the tamper material drives up the cost of the weapon, but reduces its weight and size for a given yield. As enriched uranium became relatively abundant in the US weapons program, the use of HEU (20-80% U-235) instead of natural or depleted uranium became common. A large portion of uranium produced by the US for weapons use has been in this intermediate range of enrichment. Most or all light weight strategic weapons probably use HEU tampers today.

Thorium is an inferior material for producing energy through fission compared even to depleted uranium. Its fast fission cross section is lower, so a smaller fission yield is obtained. It has a very high atomic number (Z=90), second only to uranium among practical tamper materials, and is thus a very good tamper. Its lower density may be somewhat of a disadvantage, but only because it increases the required tamper thickness - which may be undesirable in highly-volume or shape constrained designs.

A nation "breaking out" of the status as a non-weapons state might possibly consider salting a U-238 tamper with reactor-grade plutonium to increase yield. The neutron background produced could cause predetonation problems in the primary, although boron-10 shielding should mitigate it.

4.4.5.4.1 "Clean" Non-Fissile Tampers

The tamper material of choice for clean (MRR or minimum residual radiation weapons) seems to have been lead or lead-bismuth alloy. Lead is a readily available, inexpensive material, and it has the second highest atomic number of any non-fissionable element available in significant quantity (Z=82). It has been used in U.S., Soviet, and British weapon designs. The use of lead-bismuth alloy (known to have been used in British designs at least) is interesting. Bismuth may have been added to improve lead's mechanical properties, but it should also be noted that bismuth (with Z=83) has *the* highest atomic number of any available non-fissionable material. It may have thus been used to enhance lead's opacity.

When irradiated with neutrons, neither lead nor bismuth produce isotopes that constitute substantial radiological hazards. The major concern is with gamma emitters, since neither beta nor alphas are major hazards unless ingested. Neither lead nor bismuth produce isotopes that emit gammas in significant amounts.

Tungsten has also been seriously considered for use as a tamper material, although it has a relatively low atomic number (Z=74). Tungsten carbide, which is more convenient to manufacture, has also been considered. Tungsten- rhenium alloy coatings may have been used as part of tampers in some devices. Tungsten is an isotope mixture, and produces a couple isotopes with significant gamma emitting properties. Radioactive tungsten isotopes were noted in the debris released during the Redwing and Hardtack I test series.

Gold (Z=79) has been used in at least one weapon design as part of the tamper (or possibly the radiation case) - the W-71 warhead for the Spartan ABM missile. The W-71 used the thermal X-ray flux as its kill mechanism, so it was important for them to escape the weapon with as little hindrance as possible. The choice of gold may have been to tailor the opacity so that the hot X-rays present at the end of the fusion burn could escape without being absorbed. Gold is a good tamper material and has been used in ICF target designs due to its opacity.

4.4.5.4.1 "Dirty" Non-Fissile Tampers

Tamper materials may also be chosen to maximize residual radiation (relative to yield) to create radiological weapons.

Tantalum was evaluated in the US as a possible tamper material with radiological hazard potential. With Z=73, and a density of 16.65, tantalum would a very effective tamper. Natural tantalum consists of 99.988% Ta-181, which can be converted to Ta-182 with a half-life of 115 days. Ta-182 is a beta emitter with decay energy of 1.807 MeV. It emits beta particles with energies in the range of 0.25-0.54 MeV, the remaining energy is emitted as gammas (mostly in the strong gamma range of 1.19-1.22 MeV).

The most famous radiological material is cobalt. Natural cobalt consists 100% of Co-59, which becomes the energetic gamma emitter Co-60 with a 5.26 year half-life when it captures one neutron. Incorporating cobalt into a tamper could thus create an effective source of long term contamination. Cobalt has an atomic number of only 27, making it a rather poor pusher/tamper since it completely ionizes at 9.9 KeV. It could be used though if it was combined with a high Z material, and surrounded by a high-Z pusher material. No attempts to actually use cobalt in this way are known.

Zinc (Z=30)is a distant runner-up as a long-term radiological contaminant. The isotope Zn-64, which makes up 48.9% of natural zinc, would be converted to Zn-65 which is a gamma emitter with a 244 day half-life. The advantages of Zn-64 is that its faster decay leads to greater initial intensity compared to cobalt. Disadvantages are that since it makes up only half of natural zinc, it must either be isotopically enriched or the yield will be cut in half; and that it is a weaker gamma emitter than Co-60, putting out only one-fourth as many gammas for the same molar quantity. Assuming pure Zn-64 is used, the radiation intensity of Zn-65 would initially be twice as much as Co-60. This would decline to being equal in 8 months, in 5 years Co- 60 would be 110 times as intense.

Gold: I mention gold here because although it was not used in the W-71 for this purpose, it does have potential as a tailored radiological hazard. Gold consists 100% of Au-197 which breeds Au-198. Au-198 has a half-life of 2.697 days and emits 0.412 MeV gammas (along with 0.961 MeV betas). The short half-life translates into gamma emissions that are initially very intense and last for several days, but decay to low levels after several weeks. It has been considered for a short-acting battlefield or strategic radiation weapon for this reason.

4.5 Thermonuclear Weapon Designs

Since the various design elements of a thermonuclear weapon combine to form a complex integrated system, discussing the design space of these weapons involves complicated tradeoffs between design objectives and has many possible design variations.

In an attempt to address this in some kind of orderly fashion I first sketch out several basic structures for the overall weapon, in rough order of increasing sophistication (Subsection 4.5.1 Principle Design Types). Following this, I address a series of possible tradeoffs and the issues connected wit each.

4.5.1 Principle Design Types The descriptions of weapon designs, and the developmental sequence described is speculative, but it is consistent with all facts about weapons, weapon development programs, and physics of which I am currently aware.

4.5.1.1 Early Designs The earliest radiation implosion designs seem to have used a single large cylindrical chamber encompassing both the primary and cylindrical secondary. The casing was hemispherical at one end, where the primary sphere was located. The thermonuclear weapon was integral to the bomb casing itself - i.e. the ballistic shell of the bomb was the support structure for the radiation case, and the physical structure that held the entire thermonuclear device together.

Both the US and UK initially used casings made of steel, which were lined with lead or lead bismuth alloy to form the radiation case (probably 1-3 cm thick). The secondary pusher, which made up the inner wall of the radiation channel, was made of either natural uranium or lead (possibly as a lead- bismuth alloy). Operational bombs probably all used uranium tampers to maximize yield, but some test devices were equipped with lead tampers to hold down yield and fallout production. A massive radiation shield (uranium or lead) was located between the primary and secondary to prevent fuel preheating by the thermal radiation flux. A boron neutron shield was used in some designs to reduce neutron preheating.

The secondary stage consisted of the exterior pusher/tamper, a standoff gap, and a cylinder filled with fusion fuel. Lithium deuteride, highly enriched in Li-6, was the preferred fuel for maximum yield but early shortages in lithium enrichment capacity lead to the deployment of bombs containing partially enriched lithium (40% and 60% Li-6 in the U.S.), or natural lithium. Down the axis of the fusion fuel cylinder was a solid (or nearly solid) rod of plutonium or HEU for the spark plug.

The design approach of these early bombs followed that of Mike and the test devices exploded during Castle: the use of a standoff gap to create the necessary gradual compression required a large diameter (Mike was 80 inches wide, all of the Castle series devices had diameters from 54 to 61.5 inches). The rapid energy release from the primary followed by a relatively lengthy implosion required a thick casing for radiation containment, making the entire bomb very heavy. Mike weighed an anomalous 164,000 pounds, but even the Castle devices all weighed in between 23,500 and 40,000 lb.

These early bombs were thus quite massive, and had high yields. The Mk 17 and Mk 24 (the weaponized version of Castle Romeo, using unenriched lithium deuteride) had a diameter of 61.4 inches and a weight of 42,000 lb (yield: 15-20 megatons). The relatively compact and light Mk 15, whose development was completed somewhat later (and used 95% Li-6 deuteride), still had a diameter of 34.6 inches and weighed over 7,000 lb (yield: 3.8 megatons). And all of these weapons *were* bombs, since no missile could carry them. In fact, only the very largest aircraft could carry them - one per plane.

Although the primaries used in these bombs were much improved over early fission designs, they were still relatively massive initially. The TX-5 primary used in the Mike device still weighed

in at well over 1000 kg, and the comparatively thick tamper and explosive layers delayed the escape of both photons and neutrons significantly, by up to 100 nanoseconds.

4.5.1.2 Modular Weapons

During the fifties the diameter of the bomb casing and the primary shrank as US and Soviet weapons became more compact, partly driven by improved primary designs. Lighter weight megaton-range weapons were desired for greater flexibility in the types of aircraft that could carry them, and for increased payload. Light weight high yield weapons were especially important for the early ICBMs, which had limited payloads, and low accuracy. Only a light weight, high yield weapon would give a reasonable chance of destroying a designated target when carried by an ICBM. It was also useful if the same basic weapon design could be used in different weapon systems (bombs, ballistic missiles, cruise missiles, etc.)

This led to a modular approach to the weapon system. Instead of the aerodynamic casing of the delivered munition, the electronics, and the "physics package" being a single integrated entity - these three things were separated. The nuclear warhead proper (the "physics package") was self- contained, except for a cable connector to the electronics that detonated the explosives, and fired the neutron generator. The electronics package was separate, and could be different for each type of weapon (especially important for the varying fuzing requirements). These two components could then be fitted into different bomb or missile bodies to create multiple types of deployable systems.

Since the warhead casing no longer needed to withstand the environmental rigors of the completed weapon, it could be made out of lighter and less rugged materials. This led to the use of a light casing (aluminum alloy, or even plastic) that was lined with a high-Z material to form a radiation case.

4.5.1.3 Compact Light Weight Designs More efficient implosion systems and the advent of boosting made primaries more compact and less massive without sacrificing yield of efficiency. At this point (which occurred in the U.S. around 1955-1956), there seem to have been different development paths available.

One path followed the existing design principles, harnessing the increased temperatures and pressures generated by boosted light weight primaries through greater radiation confinement by increasing the thickness of the radiation case at the primary end. This evolved into a separate radiation case for the primary, a spherical shell of uranium (for example) surrounding the high explosive shell of the implosion system, with an aperture for releasing the radiation into the secondary radiation chamber (the chamber made by lining the external casing). The energy absorbed by the primary case wall at a high temperature was reradiated as the temperature in the chamber dropped. This made confinement and channeling of the thermal radiation more effective. Baffles or other barriers could be added to modulate the energy transfer into the secondary radiation case.

It appears that an alternate path may have been followed by the US starting with the Hardtack I test series (although possibly first pioneered in Redwing). According to statements made by LLNL scientists Wood and Nuckolls, and LASL Director Bradbury, new design ideas were introduced at this time that extended the Teller-Ulam concept. This coincides with the development of the very light W-47 warhead for the Polaris missile (600 lb weight and 600 kt yield, later increased to 800 kt). I speculate that the design approach introduced here was the use of modulated primary energy release.

4.5.1.4 Two Chamber Designs

At some point, the development trend toward a separate radiation case around the primary lead to a full two chamber design for the weapon, with some means of regulating radiation flow between the chambers (like a temporary radiation barrier). With better control over the radiation flux around the secondary, a reduced standoff with a reduced secondary diameter (and perhaps a lighter pusher/tamper) became possible.

This could also be conveniently combined with a spherical secondary design. This has been described as the "peanut design" - two spherical hollow chambers joined at the waist, with a primary sphere in one, and a spherical sphere in the other. Alternatively, a two chamber - spherical secondary design can be used with a modulated primary.

This approach offers the inherent advantages of spherical implosion - a smaller radius change for compression in 3-dimensions to attain a given density compared to two. Smaller radius change translates directly into faster implosion, an important consideration in a smaller, lighter, higher pressure weapon design which would be prone to disassemble faster.

In a spherical secondary the radiation shield between the primary and secondary would evolve into a baffle between the two chamber to prevent the primary from directly (and thus unevenly) heating the side of the secondary facing it, forcing the radiation flux to diffuse into the channel around the secondary.

The primary in a two-chamber design may be effectively encased in a heavy, close fitting uranium shell that can act as an implosion tamper. By trapping the explosive gases, this shell can act as the wall of a spherical piston, forcing the expanding gases to transfer all of their energy to the inward moving beryllium/plutonium shell, and minimizing the amount of explosive required. Such a primary may use a thin uranium or tungsten tamper between the beryllium and plutonium shell layers to enhance inertial confinement of the fissile mass.

4.5.1.5 Hollow Shell Designs

It was pointed out earlier that it is difficult to efficiently compress more than the outermost layers of a solid cylindrical or spherical fuel mass. In any case, only the outermost layers actually *need* to be compressed, since they contain the lion's share of the fuel mass. It would be logical then to dispense with the idea of using a solid fuel mass in the center, and only use a hollow shell of fuel in the first place. A hollow spark plug shell could be nested directly inside the fuel shell, but a second tamper layer may be included between the two.

A hollow shell could be used with either a cylindrical secondary (making it "totally tubular"), or with a spherical design.

Several advantages are obtained with this approach.

The fuel near the center that would be inefficiently compressed is eliminated, improving overall fuel utilization.

The addition of the dense second tamper or spark plug on the inner side of the fuel layer can also directly enhance compression. Whenever a shock reaches the inner side of the fuel, it will be reflected back into the fuel at higher pressure, compressing the fuel further. If the compression gradient is continuous, it will tend to "pile up" at the inner interface, with the same effect of compression enhancement. The dense inertial tamper on the inner side of the fuel layer will also help keep it at a constant high density.

Finally, the hollow shell design allows the spark plug to accelerate to very high velocities before it goes critical. The implosion velocity at criticality could be even higher than the average maximum implosion velocity for the secondary, due to the effects of thick shell collapse and convergence. An implosion velocity exceeding 1000 km/sec is conceivable. This is so fast that densities much higher than those achieved by high explosive systems would be attained before energy production from fission becomes high enough to halt implosion. Even relatively small masses of fissile material (< 1 kg) could be fissioned efficiently.

Hollow shell secondaries would be essential for use with primaries that rely on modulated energy release to create efficient compression.

4.5.1.4 High Yield and Multiple Staged Designs

The first thermonuclear devices were high yield by most any standard (10.4 Mt for Ivy Mike, 15 Mt for Castle Bravo). But they were also very heavy, and difficult to push to even greater yields. High yield weapons with greater yield-to-weight ratios, providing even higher yields in deliverable packages were desired.

As a rough approximation, we can say that the amount of energy required to implode a secondary is proportional to its mass, since the primary energy/secondary mass ratio defines the achievable implosion velocity. The yield of the secondary should also be roughly proportional to its mass. Thus there is a roughly proportional relationship between the primary and secondary yields, using similar design principles.

From available data (based on known trigger tests, and fizzles where only the primary fired), it appears that this range can be from 10-200, with 30- 50 being more typical ratios.

If a very large yield is desired, then we must obviously have a very large primary. Large fission primaries are expensive, heavy, and potentially dangerous (due to the large amount of fissile material present). Even in very heavy weapons, the yield of the primary is limited to no more than a few hundred kilotons, limiting total yield to a maximum of 10-20 megatons.

The high yield designs actually developed (mostly in the fifties and early sixties) seem to have used refined versions of the basic thermonuclear weapon design approach, as described above, with the addition of multiple staging to achieve even higher yields. The relatively light weight W-53 9 Mt warhead/bomb deployed by the US (still in service!), was one of the highest yield warheads the US ever deployed, and probably is a 3 stage weapon.

This is really large enough for almost any conceivable destructive use (except maybe blowing up asteroids). Nonetheless, military requirements for even larger weapons have been drafted, and in the case of the Soviet Union, actually built, tested, and deployed. At one point in the mid-fifties the US military requested a 60 megaton bomb! This military "requirement" was apparently driven by the fact that this was the highest yield device that could be delivered by existing aircraft. The Soviets eventually went on to develop a 100+ megaton design (tested in a 50 megaton configuration). To make such megaweapons, a bigger driving explosion is required to implode the main fusion stage. This has led to the design of three stage weapons, where a thermonuclear secondary is the main driving force to implode a gigantic tertiary stage.

Building gargantuan bombs is not the only motivation for adopting three stage weapons however. If the fusion neutrons are not harnessed to cause fission in the tamper (either because the bomb is intended to be very clean, or very dirty) then the ratio in yields between stages is correspondingly reduced - to a range of something like 10 to 15. This limits the practical maximum yield to 3 to 5 megatons. It may be doubted whether even this is much of a limitation since out of a current arsenal of over 10000 warheads, the US only has 50 bombs with yields over 3 megatons. In the fifties however this seemed unacceptably small, so "clean" weapons were deemed to require three stage design.

Three stage design can provide other advantages though. By offering the weapon designer additional freedom in design, it may be useful even if the bomb is not especially large, clean, or dirty. For example, in optimizing a weapon to minimize weight for a given yield, a designer can consider which type of driver for the main stage is the lightest - a large fission primary or a compact two stage device. If weapon-grade fissile material is very precious, then a two-stage driver might be chosen simply to minimize the over utilization of this material.

In a three stage weapon the radiation cases for the secondary and tertiary might be kept separate initially. The primary would implode the secondary but a barrier would prevent energy from reaching the tertiary. This barrier could be designed to ablate away during the secondary implosion, so that when the secondary energy release occurred, it would have become transparent.

Alternatively it may be useful to harness a portion of the primary's energy to create an initial weak compression shock in the tertiary to enhance compression efficiency.

4.5.2 "Dirty" and "Clean" Weapons

Whether to make a fission-fusion weapon into a fission-fusion-fission weapon is one of the most basic design issues. A fission-fusion weapon uses an inert (or non-fissionable) tamper and will

obtain most of its yield from the fusion reaction directly. A fission-fusion-fission weapon will obtain at least half of its yield (and often far more) from the fusion neutron induced fission of a fissionable tamper.

The basic advantage of a fission-fusion-fission weapon is that energy is extracted from a tamper which is otherwise deadweight as far as energy production in concerned. The tamper has to be there, so a lighter weapon for a given yield (or a more powerful weapon for the same weight) can be obtained without varying any other design factors. Since it is possible to do this at virtually no added cost or other penalty, compared to an inert material like lead, by using natural or depleted uranium or thorium there is basically no reason not to do it if the designer is simply interested in making big explosions.

Fission of course produces radioactive debris - fallout. Fallout can be reduced by using a material that does not become highly radioactive when bombarded by neutrons (like lead or tungsten). This requires a heavier and more expensive weapon to produce a given yield, but is also considerably reduces the short and long term contamination associated with that yield.

This is not to say that the weapon is "clean" in any commonsense meaning of the term. Neutrons escaping the weapon can still produce biohazardous carbon-14 through nitrogen capture in the air. The primary and spark plug may still contribute 10-20% fission, which for a multi-megaton weapon may still be a megaton or more of fission. Significant contamination may also occur from the "inert" tamper radioisotopes, and even from the unburned tritium produced in the fusion stage. Reducing these contributions to the lowest possible level is the realm of "minimum residual radiation" designs discussed further below.

During the fifties interest in both the US and USSR was given to developing basic design that had both clean and dirty variants. The basic design tried to minimize the essential fission yield by using a small fission primary, and spark plug sizes carefully chosen to meet ignition requirements for each stage, without being excessive (note that although only part of the spark plug will fission to ignite the fusion stage, the essentially complete fission of the remainder by fusion neutrons is inevitable). These weapons appear to have all been three-stage weapons to allow multi-megaton yields (even in the clean version) with a relatively small primary. The dirty version might simply replace the inert tamper of the tertiary with a fissionable one to boost yield.

The three-stage Bassoon and Bassoon Prime devices tested in Redwing Zuni (27 May 1956, 3.5 Mt, 15% fission) and Redwing Tewa (20 July 1956, 5 Mt, 87% fission) are US tests of this concept. Clearly though, the second test was not simply a copy of the first with a different tamper. The fusion yield dropped from 3 Mt to 0.65 Mt, and the device weight increased from 5500 kg to 7149 kg between the two tests. The inference can be made that the tertiary in the first used a large volume of relatively expensive (but light) Li-6D in a thin tamper, which was replaced by a heavier, cheaper tertiary using less fusion fuel, but a very thick fissionable tamper to capture as many neutrons as possible.

The 50 Mt three stage Tsar Bomba (King of Bombs) tested by the Soviet Union on 30 October 1961 was the largest and cleanest bomb ever tested, with 97% of its yield coming from fusion

(fission yield approximately 1.5 Mt). Assuming a primary of 250 kt (to keep the fissile content relatively low for safety reasons), we might postulate secondary and tertiary stages of 3.5 Mt and 46 Mt respectively. This fusion stages would require 1700 kg of Li6D (at 50% fusion efficiency), and something like 250 kt of fission for reliable ignition. If the initial spark plug firings were 25% efficient, later fission would release another 750 kt - placing the total at 1.25 Mt (close enough to the claimed parameters to match within the limits of accuracy).

This was a design though for a 100-150 Mt weapon! A lead tamper was used in the tested device, which could have been replaced with U-238 for the dirty version (thankfully never tested!).

4.5.3 Maximum Yield/Weight Ratio

Except for safety, the weight of a weapon required to provide a given yield is the most important design criterion. In the years since the first nuclear weapon was exploded, far more money has been spent in building nuclear weapon delivery systems than in the weapons themselves. The high cost of delivery for what is basically a rather small package is due to the fact that nuclear delivery systems are generally intended to be used only once. Clearly this is true for missiles, but it is true for bombers as well since recovery and reuse is not part of their nuclear mission profile.

Since the cost of the delivery vehicle is much greater than the cost of the warhead, making the warhead as light as possible for the intended yield quickly came to dominate the weapon design process. this is normally expressed in terms of the yield-to-weight (YTW) ratio (kt/kg).

Naturally it is easier to get a high ratio for a larger bomb. The highest ratio for any warhead in the US arsenal is the 9 Mt Mk-53/B-53 bomb, which happens to be the oldest weapon in service (operational since 1962), but also the largest. At 4000 kg, it has a ratio of 2.25 kt/kg. The Tsar Bomba, as tested, had a ratio of 1.7 kt/kg (its weight was 30 tonnes). As *designed* it had a ratio of 3.4-5 kt/kg!

Table 4.5.3-1. Yield-to-Weight Ratios of Current US Weapons

Weapon	YTW Ratio	Yield(kt)/Weight(kg)	In Service Date
Mk-53	2.25	9000/4000	1962
W-88	1.5	475/330	
W-80	1.31	170/130	
B-83	1.10	1200/1090	
W-87	1.0	300/300	
W-78	0.96	335/350	

W-76 0.61 100/165

The much earlier W-47 warhead seems to have achieved ratios of 2.2-2.7 kt/kg. However YTW ratio is not every thing. The W-87 and W-88 are said to use reduced amounts of expensive nuclear materials (deemed important when ambitious expansion of the US nuclear arsenal was planned in the early eighties) which, coupled with the much larger payloads of the MX and Trident II missiles, may account for the reduced (but still quite respectable) YTW ratios of these warheads.

Part of optimizing the YTW ratio is careful weight management. Very light weight primaries, the use of light weight weapon cases, and multiple radiation cases are innovations to minimize weight. Since the tamper is one of the heaviest parts of the weapon, squeezing as much energy out of this is very important too.

The end of surface testing of nuclear weapons after the atmospheric test ban treaty effectively removed "cleanliness" as a significant concern for designers. Complaints about fall-out vanished, and so did the ability of the international community to monitor weapon design through fall-out analysis. The cost-effectiveness of lighter weapons put great pressure on designers to extract weight saving however they could, and it is likely that the idea of using non-fissile tampers disappeared very quickly. There is scant evidence that so-called "clean" designs were ever deployed in any quantity.

The fission yield of the tamper can be increased even further by adding slow-neutron fissionable material to it. Basically this means using enriched uranium instead of natural or depleted uranium.

Highly enriched uranium is definitely known to be used in U.S. weapons. About half of the U.S. inventory of weapons-associated HEU is less than "weapons grade" (<93.4% that is). The probable use of most or all of this uranium (generally with an enrichment of 20-80%) was in thermonuclear weapon tampers.

The W-87 Peacekeeper warhead (to be redeployed on the Minuteman-III) has a current yield of 300 kt, that can be increased to 475 kt by adding a HEU sleeve or rings to the secondary. Whether this represents an actual addition to the existing secondary, or whether it replaces an existing unenriched sleeve is not known. The W-88 Trident warhead is a closely related design, and has a current yield of 475 kt indicating that it is already equipped with this addition. The 175 kt yield difference amounts to the complete fission of 10 kg of U-235.

Now, once one considers using substantial amounts of HEU in the secondary, the question of why the fusion fuel is needed at all arises. The answer: it probably is not essential. The idea of imploding fissile material is what set Stanislaw Ulam on the path to that led eventually to thermonuclear weapons. But with the availability of large amounts of HEU, and the trend toward smaller weapon yields (compared to the multimegaton behemoths of the fifties), the Ulam's idea of using radiation implosion to create a light weight high-efficiency pure fission weapon returns

as a viable possibility. It is an interesting question whether all modern strategic nuclear weapons *are* in fact thermonuclear devices!

4.5.4 Minimum Residual Radiation (MRR or "Clean") Designs

It has been pointed out elsewhere in this FAQ that ordinary fission-fusion- fission bombs (nominally 50% fission yield) are so dirty that they merit consideration as radiological weapons. Simply using a non-fissile tamper to reduce the fission yield to 5% or so helps considerably, but certainly does not result in an especially clean weapon by itself. If minimization of fallout and other sources of residual radiation is desired then considerably more effort needs to be put into design.

Minimum residual radiation designs are especially important for "peaceful nuclear explosions" (PNEs). If a nuclear explosive is to be useful for any civilian purpose, all sources of residual radiation must be reduced to the absolute lowest levels technologically possible. This means elimination neutron activation of bomb components, of materials outside the bomb, and reducing the fissile content to the smallest possible level. It may also be desirable to minimize the use of relatively hazardous materials like plutonium.

The problems of minimizing fissile yield and eliminating neutron activation are the most important. Clearly any MRR, even a small one, must be primarily a fusion device. The "clean" devices tested in the fifties and early sixties were primarily high yield strategic three-stage systems. For most uses (even military ones) these weapons are not suitable. Developing smaller yields with a low fissile content requires considerable design sophistication - small light primaries so that the low yields still produce useful radiation fluxes and high-burnup secondary designs to give a good fusion output.

Minimizing neutron activation form the abundant fusion neutrons is a serious problem since many materials inside and outside the bomb can produce hazardous activation products. The best way of avoiding this is too prevent the neutrons from getting far from the secondary. This requires using an efficient clean neutron absorber, i.e. boron-10. Ideally this should be incorporated directly into the fuel or as a lining of the fuel capsule to prevent activation of the tamper. Boron shielding of the bomb case, and the primary may be useful also.

It may be feasible to eliminate the fissile spark plug of a MRR secondary by using a centrally located deuterium-tritium spark plug similar to the way ICF capsules are ignited. Fusion bombs unavoidably produce tritium as a by- product, which can be a nuisance in PNEs.

Despite efforts to minimize radiation releases, PNEs have largely been discredited as a cost-saving civilian technology. Generally speaking, MRR devices still produce excessive radiation levels by civilian standards making their use impractical.

MRRs may have military utility as a tactical weapon, since residual contamination is slight. Such weapons are more costly and have lower performance of course.

This leads to another reason why PNEs have lost their attractiveness - there is no way to make a PNE device unsuitable for weapons use. "Peaceful" use of nuclear explosives inherently provides opportunity to develop weapons technology. As the saying goes, "the only difference between a PNE and a bomb is the tail fins".

4.5.5 Radiological Weapon Designs This is the opposite extreme of an MRR. Earlier several tamper materials were described that could be used to tailor the radioactive contamination produced by a nuclear explosion - tantalum, cobalt, zinc, and gold. Uranium tampers produce contamination in abundance - but quite a lot of energy too. In some applications it may be desired that the ratio of contamination to explosive force be increased, or tailored to a narrower spectrum of decay times compared to fission by-products.

Practical radiological weapons must incorporate the precursor isotope directly into the secondary. This is because the high compression of the secondary allows the use of reasonable masses of precursor material. In an uncompressed state, the thickness of most materials required to capture a substantial percentage of neutrons is 10-20 cm, leading to a very massive bomb. A layer of 1 cm or less will do as well when compressed by radiation implosion.

Some radioisotopes that would be very attractive for certain applications are difficult to produce in a weapon. A case in point is sodium-24, an extremely prolific producer of energetic gammas with a half-life of 14.98 hours. This isotope produces a remarkable 5.515 MeV of decay energy, with two hard gammas per decay (2.754 MeV and 1.369 MeV) and might be desired for very short-lived radiation barriers. The most obvious precursor, natural Na- 23, has a minuscule capture cross section for neutrons in the KeV range (although it is a significant hazard from induced radioactivity in soil after low altitude nuclear detonations). The best for precursor candidate for Na-24 is probably magnesium-24 (78.70% of natural magnesium) through an n,p reaction.

4.6 Weapon System Design

4.6.1 Weapon Safety

Due to their enormous destructive power, it is extremely important to ensure that nuclear weapons cannot explode at either their full yield, or at reduced yield, unless stringent and carefully specified conditions are met.

Weapons must be resist:

- malicious tampering,
- human error,
- component or systems failure (either inside or outside the weapon), and
- accidental damage.

To meet these requirements elaborate provisions for weapon safety are required. This issue has been of major concern since the first nuclear weapons, and many of the major advances in weapon design are related to weapon safety.

Weapons are invariably designed with a series of disabling mechanisms, all of which must be successfully overridden before an explosion can occur. These include locking mechanisms requiring special keys or codes, redundant safeties that must be removed to arm the weapon, environmental sensing switches (disabling mechanisms that are overridden only when the weapon has experienced environmental conditions and stresses expected during operational employment), and sophisticated fuzing systems to detonate the device at the proper place and time. Often these multiple safety systems require cooperation by more than one person to complete weapon arming.

Scenarios that must be addressed include:

- inappropriate activation of the weapon's firing system, and
- detonation of the high explosives by means other than the firing system (e.g. physical damage through fire or impact).

5.0 Effects of Nuclear Explosions

Nuclear explosions produce both immediate and delayed destructive effects. Immediate effects (blast, thermal radiation, prompt ionizing radiation) are produced and cause significant destruction within seconds or minutes of a nuclear detonation. The delayed effects (radioactive fallout and other possible environmental effects) inflict damage over an extended period ranging from hours to centuries, and can cause adverse effects in locations very distant from the site of the detonation. These two classes of effects are treated in separate subsections.

The distribution of energy released in the first minute after detonation among the three damage causing effects is:

	Low Yield (<100 kt)	High Yield (>1 Mt)
Thermal Radiation	35%	45%
Blast Wave	60%	50%
Ionizing Radiation	5%	5%
(80% gamma, 20% neutrons)		

The radioactive decay of fallout releases an additional 5-10% over time.

- 5.1 Overview of Immediate Effects
- 5.2 Overview of Delayed Effects
- 5.3 Physics of Nuclear Weapon Effects
- 5.4 Air Bursts and Surface Bursts
- 5.5 Electromagnetic Effects
- 5.6 Mechanisms of Damage and Injury

5.1 Overview of Immediate Effects

The three categories of immediate effects are: blast, thermal radiation (heat), and prompt ionizing or nuclear radiation. Their relative importance varies with the yield of the bomb. At low yields, all three can be significant sources of injury. With an explosive yield of about 2.5 kt, the three effects are roughly equal. All are capable of inflicting fatal injuries at a range of 1 km.

The equations below provide approximate scaling laws for relating the destructive radius of each effect with yield:

r_thermal = Y^0.41 * constant_th

r_blast = Y^0.33 * constant_bl

r_radiation = Y^0.19 * constant_rad

If Y is in multiples (or fractions) of 2.5 kt, then the result is in km (and all the constants equal one). This is based on thermal radiation just sufficient to cause 3rd degree burns (8 calories/cm^2); a 4.6 psi blast overpressure (and optimum burst height); and a 500 rem radiation dose.

The underlying principles behind these scaling laws are easy to explain. The fraction of a bomb's yield emitted as thermal radiation, blast, and ionizing radiation are essentially constant for all yields, but the way the different forms of energy interact with air and targets vary dramatically.

Air is essentially transparent to thermal radiation. The thermal radiation affects exposed surfaces, producing damage by rapid heating. A bomb that is 100 times larger can produce equal thermal radiation intensities over areas 100 times larger. The area of an (imaginary) sphere centered on the explosion increases with the square of the radius. Thus the destructive radius increases with the square root of the yield (this is the familiar inverse square law of electromagnetic radiation). Actually the rate of increase is somewhat less, partly due to the fact that larger bombs emit heat more slowly which reduces the damage produced by each calorie of

heat. It is important to note that the area subjected to damage by thermal radiation increases almost linearly with yield.

Blast effect is a volume effect. The blast wave deposits energy in the material it passes through, including air. When the blast wave passes through solid material, the energy left behind causes damage. When it passes through air it simply grows weaker. The more matter the energy travels through, the smaller the effect. The amount of matter increases with the volume of the imaginary sphere centered on the explosion. Blast effects thus scale with the inverse cube law which relates radius to volume.

The intensity of nuclear radiation decreases with the inverse square law like thermal radiation. However nuclear radiation is also strongly absorbed by the air it travels through, which causes the intensity to drop off much more rapidly.

These scaling laws show that the effects of thermal radiation grow rapidly with yield (relative to blast), while those of radiation rapidly decline.

In the Hiroshima attack (bomb yield approx. 15 kt) casualties (including fatalities) were seen from all three causes. Burns (including those caused by the ensuing fire storm) were the most prevalent serious injury (two thirds of those who died the first day were burned), and occurred at the greatest range. Blast and burn injuries were both found in 60-70% of all survivors. People close enough to suffer significant radiation illness were well inside the lethal effects radius for blast and flash burns, as a result only 30% of injured survivors showed radiation illness. Many of these people were sheltered from burns and blast and thus escaped their main effects. Even so, most victims with radiation illness also had blast injuries or burns as well.

With yields in the range of hundreds of kilotons or greater (typical for strategic warheads) immediate radiation injury becomes insignificant. Dangerous radiation levels only exist so close to the explosion that surviving the blast is impossible. On the other hand, fatal burns can be inflicted well beyond the range of substantial blast damage. A 20 megaton bomb can cause potentially fatal third degree burns at a range of 40 km, where the blast can do little more than break windows and cause superficial cuts.

It should be noted that the atomic bombings of Hiroshima and Nagasaki caused fatality rates were ONE TO TWO ORDERS OF MAGNITUDE higher than the rates in conventional fire raids on other Japanese cities. Eventually on the order of 200,000 fatalities, which is about one-quarter of all Japanese bombing deaths, occurred in these two cities with a combined population of less than 500,000. This is due to the fact that the bombs inflicted damage on people and buildings virtually instantaneously and without warning, and did so with the combined effects of flash, blast, and radiation. Widespread fatal injuries were thus inflicted instantly, and the many more people were incapacitated and thus unable to escape the rapidly developing fires in the suddenly ruined cities. Fire raids in comparison, inflicted few immediate or direct casualties; and a couple of hours elapsed from the raid's beginning to the time when conflagrations became general, during which time the population could flee.

A convenient rule of thumb for estimating the short-term fatalities from all causes due to a nuclear attack is to count everyone inside the 5 psi blast overpressure contour around the hypocenter as a fatality. In reality, substantial numbers of people inside the contour will survive and substantial numbers outside the contour will die, but the assumption is that these two groups will be roughly equal in size and balance out. This completely ignores any possible fallout effects.

5.2 Overview of Delayed Effects

5.2.1 Radioactive Contamination

The chief delayed effect is the creation of huge amounts of radioactive material with long lifetimes (half-lifes ranging from days to millennia). The primary source of these products is the debris left from fission reactions. A potentially significant secondary source is neutron capture by non-radioactive isotopes both within the bomb and in the outside environment.

When atoms fission they can split in some 40 different ways, producing a mix of about 80 different isotopes. These isotopes vary widely in stability, some our completely stable while others undergo radioactive decay with half-lifes of fractions of a second. The decaying isotopes may themselves form stable or unstable daughter isotopes. The mixture thus quickly becomes even more complex, some 300 different isotopes of 36 elements have been identified in fission products.

Short-lived isotopes release their decay energy rapidly, creating intense radiation fields that also decline quickly. Long-lived isotopes release energy over long periods of time, creating radiation that is much less intense but more persistent. Fission products thus initially have a very high level of radiation that declines quickly, but as the intensity of radiation drops, so does the rate of decline.

A useful rule-of-thumb is the "rule of sevens". This rule states that for every seven-fold increase in time following a fission detonation (starting at or after 1 hour), the radiation intensity decreases by a factor of 10. Thus after 7 hours, the residual fission radioactivity declines 90%, to one-tenth its level of 1 hour. After 7*7 hours (49 hours, approx. 2 days), the level drops again by 90%. After 7*2 days (2 weeks) it drops a further 90%; and so on for 14 weeks. The rule is accurate to 25% for the first two weeks, and is accurate to a factor of two for the first six months. After 6 months, the rate of decline becomes much more rapid. The rule of sevens corresponds to an approximate t^-1.2 scaling relationship.

These radioactive products are most hazardous when they settle to the ground as "fallout". The rate at which fallout settles depends very strongly on the altitude at which the explosion occurs, and to a lesser extent on the size of the explosion.

If the explosion is a true air-burst (the fireball does not touch the ground), when the vaporized radioactive products cool enough to condense and solidify, they will do so to form microscopic particles. These particles are mostly lifted high into the atmosphere by the rising fireball, although significant amounts are deposited in the lower atmosphere by mixing that occurs due to convective circulation within the fireball. The larger the explosion, the higher and faster the fallout is lofted, and the smaller the proportion that is deposited in the lower atmosphere. For explosions with yields of 100 kt or less, the fireball does not rise abve the troposphere where precipitation occurs. All of this fallout will thus be brought to the ground by weather processes within months at most (usually much faster). In the megaton range, the fireball rises so high that it enters the stratosphere. The stratosphere is dry, and no weather processes exist there to bring fallout down quickly. Small fallout particles will descend over a period of months or years. Such long-delayed fallout has lost most of its hazard by the time it comes down, and will be distributed on a global scale. As yields increase above 100 kt, progressively more and more of the total fallout is injected into the stratosphere.

An explosion closer to the ground (close enough for the fireball to touch) sucks large amounts of dirt into the fireball. The dirt usually does not vaporize, and if it does, there is so much of it that it forms large particles. The radioactive isotopes are deposited on soil particles, which can fall quickly to earth. Fallout is deposited over a time span of minutes to days, creating downwind contamination both nearby and thousands of kilometers away. The most intense radiation is created by nearby fallout, because it is more densely deposited, and because short-lived isotopes haven't decayed yet. Weather conditions can affect this considerably of course. In particular, rainfall can "rain out" fallout to create very intense localized concentrations. Both external exposure to penetrating radiation, and internal exposure (ingestion of radioactive material) pose serious health risks.

Explosions close to the ground that do not touch it can still generate substantial hazards immediately below the burst point by neutron-activation. Neutrons absorbed by the soil can generate considerable radiation for several hours.

The megaton class weapons that were developed in the US and USSR during the fifties and sixties have been largely retired, being replaced with much smaller yield warheads. The yield of a modern strategic warhead is, with few exceptions, now typically in the range of 200-750 kt. Recent work with sophisticated climate models has shown that this reduction in yield results in a much larger proportion of the fallout being deposited in the lower atmosphere, and a much faster and more intense deposition of fallout than had been assumed in studies made during the sixties and seventies. The reduction in aggregate strategic arsenal yield that occurred when high yield weapons were retired in favor of more numerous lower yield weapons has actually increased the fallout risk.

5.2.2 Effects on the Atmosphere and Climate

Although not as directly deadly as fallout, other environmental effects can be quite harmful.

5.2.2.1 Harm to the Ozone Layer

The high temperatures of the nuclear fireball, followed by rapid expansion and cooling, cause large amounts of nitrogen oxides to form from the oxygen and nitrogen in the atmosphere (very similar to what happens in combustion engines). Each megaton of yield will produce some 5000 tons of nitrogen oxides. The rising fireball of a high kiloton or megaton range warhead will carry these nitric oxides well up into the stratosphere, where they can reach the ozone layer. A series of large atmospheric explosions could significantly deplete the ozone layer. The high yield tests in the fifties and sixties probably did cause significant depletion, but the ozone measurements made at the time were too limited to pick up the expected changes out of natural variations.

5.2.2.2 Nuclear Winter

The famous TTAPS (Turco, Toon, Ackerman, Pollack, and Sagan) proposal regarding a potential "nuclear winter" is another possible occurrence. This effect is caused by the absorption of sunlight when large amounts of soot are injected into the atmosphere by the widespread burning of cities and petroleum stocks destroyed in a nuclear attack.

Similar events have been observed naturally when large volcanic eruptions have injected large amounts of dust into the atmosphere. The Tambora eruption of 1815 (the largest volcanic eruption in recent history) was followed by "the year without summer" in 1816, the coldest year in the last few centuries.

Soot is far more efficient in absorbing light than volcanic dust, and soot particles are small and hydrophobic and thus tend not to settle or wash out as easily.

Although the initial TTAPS study was met with significant skepticism and criticism, later and more sophisticated work by researchers around the world have confirmed it in all essential details. These studies predict that the amount of soot that would be produced by burning most of the major cities in the US and USSR would severly disrupt climate on a world-wide basis. The major effect would be a rapid and drastic reduction in global temperature, especially over land. All recent studies indicate that if large scale nucelar attack occur against urban or petrochemical targets, average temperature reductions of at least 10 degrees C would occur lasting many months. This level of cooling far exceeds any that has been observed in recorded history, and is comparable to that of a full scale ice age. In areas downwind from attack sites, the cooling can reach 35 degrees C. It is probable that no large scale temperature excursion of this size has occurred in 65 million years.

Smaller attacks would create reduced effects of course. But it has been pointed out that most of the world's food crops are subtropical plants that would have dramatic drops in productivity if an average temperature drop of even one degree were to occur for even a short time during the growing season. Since the world maintains a stored food supply equal to only a few months of consumption, a war during the Northern Hemisphere spring or summer could still cause deadly starvation around the globe from this effect alone even if it only produced a mild "nuclear autumn".

5.3 Physics of Nuclear Weapon Effects

Thermal radiation and blast are inevitable consequences of the near instantaneous release of an immense amount of energy in a very small volume, and are thus characteristic to all nuclear weapons regardless of type or design details. The release of ionizing radiation, both at the instant of explosion and delayed radiation from fallout, is governed by the physics of the nuclear reactions involved and how the weapon is constructed, and is thus very dependent on both weapon type and design.

5.3.1 Fireball Physics

The fireball is the hot ball of gas created when a nuclear explosion heats the bomb itself, and the immediate surrounding environment, to very high temperatures. As this incandescent ball of hot gas expands, it radiates part of its energy away as thermal radiation (including visible and ultraviolet light), part of its energy also goes into creating a shock wave or blast wave in the surrounding environment. The generation of these two destructive effects are thus closely linked by the physics of the fireball. In the discussion below I assume the fireball is forming in open air, unless stated otherwise.

5.3.1.1 The Early Fireball

Immediately after the energy-producing nuclear reactions in the weapon are completed, the energy is concentrated in the nuclear fuels themselves. The energy is stored as (in order of importance): thermal radiation or photons; as kinetic energy of the ionized atoms and the electrons (mostly as electron kinetic energy since free electrons outnumber the atoms); and as excited atoms, which are partially or completely stripped of electrons (partially for heavy elements, completely for light ones).

Thermal (also called blackbody) radiation is emitted by all matter. The intensity and most prevalent wavelength is a function of the temperature, both increasing as temperature increases. The intensity of thermal radiation increases very rapidly - as the fourth power of the temperature. Thus at the 60-100 million degrees C of a nuclear explosion, which is some 10,000 times hotter than the surface of the sun, the brightness (per unit area) is some 10 quadrillion (10^16) times greater! Consequently about 80% of the energy in a nuclear explosion exists as photons. At these temperatures the photons are soft x-rays with energies in the range of 10-200 KeV.

The first energy to escape from the bomb are the gamma rays produced by the nuclear reactions. They have energies in the MeV range, and a significant number of them penetrate through the tampers and bomb casing and escape into the outside world at the speed of light. The gamma rays strike and ionize the surrounding air molecules, causing chemical reactions that form a dense layer of "smog" tens of meters deep around the bomb. This smog is composed primarily of ozone, and nitric and nitrous oxides.

X-rays, particularly the ones at the upper end of the energy range, have substantial penetrating power and can travel significant distances through matter at the speed of light before being absorbed. Atoms become excited when they absorb x-rays, and after a time they re-emit part of the energy as a new lower energy x-ray. By a chain of emissions and absorptions, the x-rays carry energy out of the hot center of the bomb, a process called radiative transport. Since each absorption/re-emission event takes a certain amount of time, and the direction of re-emission is random (as likely back toward the center of the bomb as away from it), the net rate of radiative transport is considerably slower than the speed of light. It is however initially much faster than the expansion of the plasma (ionized gas) making up the fireball or the velocity of the neutrons.

An expanding bubble of very high temperatures is thus formed called the "iso-thermal sphere". It is a sphere were everything has been heated by x-rays to a nearly uniform temperature, initially in the tens of millions of degrees. As soon as the sphere expands beyond the bomb casing it begins radiating light away through the air (unless the bomb is buried or underwater). Due to the still enormous temperatures, it is incredibly brilliant (surface brightness trillions of times more intense than the sun). Most of the energy being radiated is in the x-ray and far ultraviolet range to which air is not transparent. Even at the wavelengths of the near ultraviolet and visible light, the "smog" layer absorbs much of the energy. Then too, at this stage the fireball is only a few meters across. Thus the apparent surface brightness at a distance, and the output power (total brightness) is not nearly as intense as the fourth-power law would indicate.

5.3.1.2 Blast Wave Development and Thermal Radiation Emission

As the fireball expands, it cools and the wavelength of the photons transporting energy drops. Longer wavelength photons do not penetrate as far before being absorbed, so the speed of energy transport also drops. When the isothermal sphere cools to about 300,000 degrees C (and the surface brightness has dropped to being a mere 10 million times brighter than the sun), the rate of radiative growth is about equal to the speed of sound in the fireball plasma. At this point a shock wave forms at the surface of the fireball as the kinetic energy of the fast moving ions starts transferring energy to the surrounding air. This phenomenon, known as "hydrodynamic separation", occurs for a 20 kt explosion about 100 microseconds after the explosion, when the fireball is some 13 meters across. A shock wave internal to the fireball caused by the rapidly expanding bomb debris may overtake and reinforce the fireball surface shock wave a few hundred microseconds later.

The shock wave initially moves at some 30 km/sec, a hundred times the speed of sound in normal air. This compresses and heats the air enormously, up to 30,000 degrees C (some five times the sun's surface temperature). At this temperature the air becomes ionized and incandescent. Ionized gas is opaque to visible radiation, so the glowing shell created by the shock front hides the much hotter isothermal sphere inside. The shock front is many times brighter than the sun, but since it is much dimmer than the isothermal sphere it acts as an optical shutter, causing the fireball's thermal power to drop rapidly.

The fireball is at its most brilliant just as hydrodynamic separation occurs, the great intensity compensating for the small size of the fireball. The rapid drop in temperature causes the thermal

power to drop ten-fold, reaching a minimum in about 10 milliseconds for a 20 kt bomb (100 milliseconds for 1 Mt bomb). This "first pulse" contains only about 1 percent of the bomb's total emitted thermal radiation. At this minimum, the fireball of a 20 kt bomb is 180 meters across.

As the shock wave expands and cools to around 3000 degrees, it stops glowing and gradually also becomes transparent. This is called "breakaway" and occurs at about 15 milliseconds for a 20 kt bomb, when the shock front has expanded to 220 meters and is travelling at 4 km/second. The isothermal sphere, at a still very luminous 8000 degrees, now becomes visible and both the apparent surface temperature and brightness of the fireball climb to form the "second pulse". The isothermal sphere has grown considerably in size and now consists almost entirely of light at wavelengths to which air is transparent, so it regains much of the total luminosity of the first peak despite its lower temperature. This second peak occurs at 150 milliseconds for a 20 kt bomb, at 900 milliseconds for a 1 Mt bomb. After breakaway, the shock (blast) wave and the fireball do not interact further.

A firm cutoff for this second pulse is impossible to provide because the emission rate gradually declines over an extended period. Some rough guidelines are that by 300 milliseconds for a 20 kt bomb (1.8 seconds for a 1 Mt) 50% of the total thermal radiation has been emitted, and the rate has dropped to 40% of the second peak. These figures become 75% total emitted and 10% peak rate by 750 milliseconds (20 kt) and 4.5 second (1 Mt). The emission time scales roughly as the 0.45 power of yield (Y^0.45).

Although this pulse never gets as bright as the first, it emits about 99% of the thermal radiation because it is so much longer.

5.3.2 Ionizing Radiation Physics

There are four types of ionizing radiation produced by nuclear explosions that can cause significant injury: neutrons, gamma rays, beta particles, and alpha particles. Gamma rays are energetic (short wavelength) photons (as are X-rays), beta particles are energetic (fast moving) electrons, and alpha particles are energetic helium nuclei. Neutrons are damaging whether they are energetic or not, although the faster they are, the worse their effects.

They all share the same basic mechanism for causing injury though: the creation of chemically reactive compounds called "free radicals" that disrupt the normal chemistry of living cells. These radicals are produced when the energetic radiation strikes a molecule in the living issue, and breaks it into ionized (electrically charged) fragments. Fast neutrons can do this also, but all neutrons can also transmute ordinary atoms into radioactive isotopes, creating even more ionizing radiation in the body.

The different types of radiation present different risks however. Neutrons and gamma rays are very penetrating types of radiation. They are the hardest to stop with shielding. They can travel through hundreds of meters of air and the walls of ordinary houses. They can thus deliver deadly radiation doses even if an organism is not in immediate contact with the source. Beta particles are less penetrating, they can travel through several meters of air, but not walls, and

can cause serious injury to organisms that are near to the source. Alpha particles have a range of only a few centimeters in air, and cannot even penetrate skin. Alphas can only cause injury if the emitting isotope is ingested.

The shielding effect of various materials to radiation is usually expressed in half-value thickness, or tenth-value thickness: in other words, the thickness of material required to reduce the intensity of radiation by one-half or one-tenth. Successive layers of shielding each reduce the intensity by the same proportion, so three tenth-value thickness reduce the intensity to one-thousandth (a tenth-value thickness is about 3.3 half-value thicknesses). Some example tenth-value thicknesses for gamma rays are: steel 8.4-11 cm, concrete 28-41 cm, earth 41-61 cm, water 61-100 cm, and wood 100-160 cm. The thickness ranges indicate the varying shielding effect for different gamma ray energies.

Even light clothing provides substantial shielding to beta rays.

5.3.2.1 Sources of Radiation

5.3.2.1.1 Prompt Radiation

Radiation is produced directly by the nuclear reactions that generate the explosion, and by the decay of radioactive products left over (either fission debris, or induced radioactivity from captured neutrons).

The explosion itself emits a very brief burst (about 100 nanoseconds) of gamma rays and neutrons, before the bomb has blown itself apart. The intensity of these emissions depends very heavily on the type of weapon and the specific design. In most designs the initial gamma ray burst is almost entirely absorbed by the bomb (tamper, casing, explosives, etc.) so it contributes little to the radiation hazard. The neutrons, being more penetrating, may escape. Both fission and fusion reactions produce neutrons. Fusion produces many more of them per kiloton of yield, and they are generally more energetic than fission neutrons. Some weapons (neutron bombs) are designed specifically to emit as much energy in the form as neutrons as possible. In heavily tamped fission bombs few if any neutrons escape. It is estimated that no significant neutron exposure occurred from Fat Man, and only 2% of the total radiation dose from Little Boy was due to neutrons.

The neutron burst itself can be a significant source of radiation, depending on weapon design. As the neutrons travel through the air they are slowed by collisions with air atoms, and are eventually captured. Even this process of neutron attenuation generates hazardous radiation. Part of the kinetic energy lost by fast neutrons as they slow is converted into gamma rays, some with very high energies (for the 14.1 MeV fusion neutrons). The duration of production for these neutron scattering gammas is about 10 microseconds. The capture of neutrons by nitrogen-14 also produces gammas, a process completed by 100 milliseconds.

Immediately after the explosion, there are substantial amounts of fission products with very short half-lifes (milliseconds to minutes). The decay of these isotopes generate correspondingly

intense gamma radiation that is emitted directly from the fireball. This process is essentially complete within 10 seconds.

The relative importance of these gamma ray sources depends on the size of the explosion. Small explosions (20 kt, say) can generate up to 25% of the gamma dose from the direct gammas and neutron reactions. For large explosions (1 Mt) this contribution is essentially zero. In all cases, the bulk of the gammas are produced by the rapid decay of radioactive debris.

5.3.2.1.2 Delayed Radiation

Radioactive decay is the sole source of beta and alpha particles. They are also emitted during the immediate decay mentioned above of course, but their range is too short to make any prompt radiation contribution. Betas and alphas become important when fallout begins settling out. Gammas remain very important at this stage as well.

Fallout is a complex mixture of different radioactive isotopes, the composition of which continually changes as each isotope decays into other isotopes. Many isotopes make significant contributions to the overall radiation level. Radiation from short lived isotopes dominates initially, and the general trend is for the intensity to continually decline as they disappear. Over time the longer lived isotopes become increasingly important, and a small number of isotopes emerge as particular long-term hazards.

Radioactive isotopes are usually measured in terms of curies. A curie is the quantity of radioactive material that undergoes 3.7x10^10 decays/sec (equal to 1 g of radium-226). More recently the SI unit bequerel has become common in scientific literature, one bequerel is 1 decay/sec . The fission of 57 grams of material produces 3x10^23 atoms of fission products (two for each atom of fissionable material). One minute after the explosion this mass is undergoing decays at a rate of 10^21 disintegrations/sec (3x10^10 curies). It is estimated that if these products were spread over 1 km^2, then at a height of 1 m above the ground one hour after the explosion the radiation intensity would be 7500 rads/hr.

Isotopes of special importance include iodine-131, strontium-90 and 89, and cesium-137. This is due to both their relative abundance in fallout, and to their special biological affinity. Isotopes that are readily absorbed by the body, and concentrated and stored in particular tissues can cause harm out of proportion to their abundance.

Iodine-131 is a beta and gamma emitter with a half-life of 8.07 days (specific activity 124,000 curies/g) Its decay energy is 970 KeV; usually divided between 606 KeV beta, 364 KeV gamma. Due to its short half-life it is most dangerous in the weeks immediately after the explosion, but hazardous amounts can persist for a few months. It constitutes some 2% of fission-produced isotopes - 1.6x10^5 curies/kt. Iodine is readily absorbed by the body and concentrated in one small gland, the thyroid.

Strontium-90 is a beta emitter (546 KeV, no gammas) with a half-life of 28.1 years (specific activity 141 curies/g), Sr-89 is a beta emitter (1.463 MeV, gammas very rarely) with a half-life of 52 days (specific activity 28,200 Ci/g). Each of these isotopes constitutes about 3% of total

fission isotopes: 190 curies of Sr-90 and 3.8x10^4 curies of Sr-89 per kiloton. Due to their chemical resemblance to calcium these isotopes are absorbed fairly well, and stored in bones. Sr-89 is an important hazard for a year or two after an explosion, but Sr-90 remains a hazard for centuries. Actually most of the injury from Sr-90 is due to its daughter isotope yttrium-90. Y-90 has a half-life of only 64.2 hours, so it decays as fast as it is formed, and emits 2.27 MeV beta particles.

Cesium-137 is a beta and gamma emitter with a half-life of 30.0 years (specific activity 87 Ci/g). Its decay energy is 1.176 MeV; usually divided by 514 KeV beta, 662 KeV gamma. It comprises some 3-3.5% of total fission products - 200 curies/kt. It is the primary long-term gamma emitter hazard from fallout, and remains a hazard for centuries.

Although not important for acute radiation effects, the isotopes carbon-14 and tritium are also of interest because of possible genetic injury. These are not direct fission products. They are produced by the interaction of fission and fusion neutrons with the atmosphere and, in the case of tritium, as a direct product of fusion reactions. Most of the tritium generated by fusion is consumed in the explosion but significant amounts survive. Tritium is also formed by the capture of fast neutrons by nitrogen atoms in the air: N-14 + n -> T + C-12. Carbon-14 in also formed by neutron-nitrogen reactions: N-14 + n -> C-14 + p. Tritium is a very weak beta emitter (18.6 KeV, no gamma) with a half-life of 12.3 years (9700 Ci/g).

Carbon-14 is also a weak beta emitter (156 KeV, no gamma), with a half-life of 5730 years (4.46 Ci/g). Atmospheric testing during the fifties and early sixties produced about 3.4 g of C-14 per kiloton (15.2 curies) for a total release of 1.75 tonnes (7.75x10^6 curies). For comparison, only about 1.2 tonnes of C-14 naturally exists, divided between the atmosphere (1 tonne) and living matter (0.2 tonne). Another 50-80 tonnes is dissolved in the oceans. Due to carbon exchange between the atmosphere and oceans, the half-life of C-14 residing in the atmosphere is only about 6 years. By now the atmospheric concentration has returned to within 1% or so of normal. High levels of C-14 remain in organic material formed during the sixties (in wood, say, or DNA).

5.4 Air Bursts and Surface Bursts

It might seem logical that the most destructive way of using a nuclear weapon would be to explode it right in the middle of its target - i.e. ground level. But for most uses this is not true. Generally nuclear weapons are designed to explode above the ground - as air bursts (the point directly below the burst point is called the hypocenter). Surface (and sub-surface) bursts can be used for special purposes.

5.4.1 Air Bursts

When an explosion occurs it sends out a shock wave like an expanding soap bubble. If the explosion occurs above the ground the bubble expands and when it reaches the ground it is reflected - i.e. the shock front bounces off the ground to form a second shock wave travelling

behind the first. This second shock wave travels faster than the first, or direct, shock wave since it is travelling through air already moving at high speed due to the passage of the direct wave. The reflected shock wave tends to overtake the direct shock wave and when it does they combine to form a single reinforced wave.

This is called the Mach Effect, and produces a skirt around the base of the shock wave bubble where the two shock waves have combined. This skirt sweeps outward as an expanding circle along the ground with an amplified effect compared to the single shock wave produced by a ground burst.

The higher the burst altitude, the weaker the shock wave is when it first reaches the ground. On the other hand, the shock wave will also affect a larger area. Air bursts therefore reduce the peak intensity of the shock wave, but increase the area over which the blast is felt. For a given explosion yield, and a given blast pressure, there is a unique burst altitude at which the area subjected to that pressure is maximized. This is called the optimum burst height for that yield and pressure.

All targets have some level of vulnerability to blast effects. When some threshold of blast pressure is reached the target is completely destroyed. Subjecting the target to pressures higher than that accomplishes nothing. By selecting an appropriate burst height, an air burst can destroy a much larger area for most targets than can surface bursts.

The Mach Effect enhances shock waves with pressures below 50 psi. At or above this pressure the effect provides very little enhancement, so air bursts have little advantage if very high blast pressures are desired.

An additional effect of air bursts is that thermal radiation is also distributed in a more damaging fashion. Since the fireball is formed above the earth, the radiation arrives at a steeper angle and is less likely to be blocked by intervening obstacles and low altitude haze.

5.4.2 Surface Bursts

Surface bursts are useful if local fallout is desired, or if the blast is intended to destroy a buried or very hard structure like a missile silo or a dam. Shock waves are transmitted through the soil more effectively if the bomb is exploded in immediate contact with it, so ground bursts would be used for destroying buried command centers and the like. Some targets, like earth-fill dams, require actual cratering to be destroyed and would be ground burst targets.

5.4.3 Sub-Surface Bursts

Exploding a bomb below ground level can be even more effective for producing craters and destroying buried structures. It can also eliminate thermal radiation and reduce the range of blast effects substantially. The problem, of course is getting the bomb underground. Earth-penetrating bombs have been developed that can punch over one hundred feet into the earth.

5.5 Electromagnetic Effects

The high temperatures and energetic radiation produced by nuclear explosions also produce large amounts of ionized (electrically charged) matter which is present immediately after the explosion. Under the right conditions, intense currents and electromagnetic fields can be produced, generically called EMP (Electromagnetic Pulse), that are felt at long distances. Living organisms are impervious to these effects, but electrical and electronic equipment can be temporarily or permanently disabled by them. Ionized gases can also block short wavelength radio and radar signals (fireball blackout) for extended periods.

The occurrence of EMP is strongly dependent on the altitude of burst. It can be significant for surface or low altitude bursts (below 4,000 m); it is very significant for high altitude bursts (above 30,000 m); but it is not significant for altitudes between these extremes. This is because EMP is generated by the asymmetric absorption of instantaneous gamma rays produced by the explosion. At intermediate altitudes the air absorbs these rays fairly uniformly and does not generate long range electromagnetic disturbances.

The formation EMP begins with the very intense, but very short burst of gamma rays caused by the nuclear reactions in the bomb. About 0.3% of the bomb's energy is in this pulse, but it lasts for only 10 nanoseconds or so. These gamma rays collide with electrons in air molecules, and eject the electrons at high energies through a process called Compton scattering. These energetic electrons in turn knock other electrons loose, and create a cascade effect that produces some 30,000 electrons for every original gamma ray.

In low altitude explosions the electrons, being very light, move much more quickly than the ionized atoms they are removed from and diffuse away from the region where they are formed. This creates a very strong electric field which peaks in intensity at 10 nanoseconds. The gamma rays emitted downward however are absorbed by the ground which prevents charge separation from occurring. This creates a very strong vertical electric current which generates intense electromagnetic emissions over a wide frequency range (up to 100 MHZ) that emanate mostly horizontally. At the same time, the earth acts as a conductor allowing the electrons to flow back toward the burst point where the positive ions are concentrated. This produces a strong magnetic field along the ground. Although only about 3x10^-10 of the total explosion energy is radiated as EMP in a ground burst (10^6 joules for 1 Mt bomb), it is concentrated in a very short pulse. The charge separation persists for only a few tens of microseconds, making the emission power some 100 gigawatts. The field strengths for ground bursts are high only in the immediate vicinity of the explosion. For smaller bombs they aren't very important because they are strong only where the destruction is intense anyway. With increasing yields, they reach farther from the zone of intense destruction. With a 1 Mt bomb, they remain significant out to the 2 psi overpressure zone (5 miles).

High altitude explosions produce EMPs that are dramatically more destructive. About 3x10^-5 of the bomb's total energy goes into EMP in this case, 10^11 joules for a 1 Mt bomb. EMP is formed in high altitude explosions when the downwardly directed gamma rays encounter denser layers of air below. A pancake shaped ionization region is formed below the bomb. The zone

can extend all the way to the horizon, to 2500 km for an explosion at an altitude of 500 km. The ionization zone is up to 80 km thick at the center. The Earth's magnetic field causes the electrons in this layer to spiral as they travel, creating a powerful downward directed electromagnetic pulse lasting a few microseconds. A strong vertical electrical field (20-50 KV/m) is also generated between the Earth's surface and the ionized layer, this field lasts for several minutes until the electrons are recaptured by the air. Although the peak EMP field strengths from high altitude bursts are only 1-10% as intense as the peak ground burst fields, they are nearly constant over the entire Earth's surface under the ionized region.

The effects of these field on electronics is difficult to predict, but can be profound. Enormous induced electric currents are generated in wires, antennas, and metal objects (like missiles, airplanes, and building frames). Commercial electrical grids are immense EMP antennas and would be subjected to voltage surges far exceeding those created by lightning, and over vastly greater areas. Modern VLSI chips are extremely sensitive to voltage surges, and would be burned out by even small leakage currents. Military equipment is generally designed to be resistant to EMP, but realistic tests are very difficult to perform and EMP protection rests on attention to detail. Minor changes in design, incorrect maintenance procedures, poorly fitting parts, loose debris, moisture, and ordinary dirt can all cause elaborate EMP protections to be totally circumvented. It can be expected that a single high yield, high altitude explosion over an industrialized area would cause massive disruption for an indeterminable period, and would cause huge economic damages (all those damaged chips add up).

A separate effect is the ability of the ionized fireball to block radio and radar signals. Like EMP, this effect becomes important with high altitude bursts. Fireball blackout can cause radar to be blocked for tens of seconds to minutes over an area tens of kilometers across. High frequency radio can be disrupted over hundreds to thousands of kilometers for minutes to hours depending on exact conditions.

5.6 Mechanisms of Damage and Injury

The different mechanisms are discussed individually, but it should be no surprise that in combination they often accentuate the harm caused by each other. I will discuss such combined effects wherever appropriate.

5.6.1 Thermal Damage and Incendiary Effects

Thermal damage from nuclear explosions arises from the intense thermal (heat) radiation produced by the fireball. The thermal radiation (visible and infrared light) falls on exposed surfaces and is wholly or partly absorbed. The radiation lasts from about a tenth of a second, to several seconds depending on bomb yield (it is longer for larger bombs). During that time its intensity can exceed 1000 watts/cm^2 (the maximum intensity of direct sunlight is 0.14 watts/cm^2). For a rough comparison, the effect produced is similar to direct exposure to the flame of an acetylene torch.

The heat is absorbed by the opaque surface layer of the material on which it falls, which is usually a fraction of a millimeter thick. Naturally dark materials absorb more heat than light colored or reflective ones. The heat is absorbed much faster than it can be carried down into the material through conduction, or removed by reradiation or convection, so very high temperatures are produced in this layer almost instantly. Surface temperatures can exceed 1000 degrees C close to the fireball. Such temperatures can cause dramatic changes to the material affected, but they do not penetrate in very far.

More total energy is required to inflict a given level of damage for a larger bomb than a smaller one since the heat is emitted over a longer period of time, but this is more than compensated for by the increased thermal output. The thermal damage for a larger bomb also penetrates further due to the longer exposure.

Thermal radiation damage depends very strongly on weather conditions. Cloud cover, smoke, or other obscuring material in the air can considerably reduce effective damage ranges over clear air conditions.

For all practical purposes, the emission of thermal radiation by a bomb is complete by the time the shock wave arrives. Regardless of yield, this generalization is only violated in the area of total destruction around a nuclear explosion where 100% mortality would result from any one of the three damage effects.

Incendiary effects refer to anything that contributes to the occurrence of fires after the explosion, which is a combination of the effects of thermal radiation and blast.

5.6.1.1 Thermal Injury

The result of very intense heating of skin is to cause burn injuries. The burns caused by the sudden intense thermal radiation from the fireball are called "flash burns". The more thermal radiation absorbed, the more serious the burn. The table below indicates the amount of thermal radiation required to cause different levels of injury, and the maximum ranges at which they occur, for different yields of bombs. The unit of heat used are gram-calories, equal to 4.2 joules (4.2 watts for 1 sec). Skin color significantly affects susceptibility, light skin being less prone to burns. The table assumes medium skin color.

SEVERITY	20 Kilotons	1 Megaton	20 Megatons
1st Degree	2.5 cal/cm^2 (4.3 km)	3.2 cal/cm^2 (18 km)	5 cal/cm^2 (52 km)
2nd Degree	5 cal/cm^2 (3.2 km)	6 cal/cm^2 (14.4 km)	8.5 cal/cm^2 (45 km)
3rd Degree	8 cal/cm^2 (2.7 km)	10 cal/cm^2 (12 km)	12 cal/cm^2 (39 km)

Convenient scaling laws to allow calculation of burn effects for any yield are:

r_thermal_1st = Y^0.38 * 1.20

r_thermal_2nd = Y^0.40 * 0.87

r_thermal_3rd = Y^0.41 * 0.67

Range is in km, yield is in kt; the equations are accurate to within 10% or so from 1 kt to 20 Mt.

First degree flash burns are not serious, no tissue destruction occurs. They are characterized by immediate pain, followed by reddening of the skin. Pain and sensitivity continues for some minutes or hours, after which the affected skin returns to normal without further incident.

Second degree burns cause damage to the underlying dermal tissue, killing some portion of it. Pain and redness is followed by blistering within a few hours as fluids collect between the epidermis and damaged tissue. Sufficient tissue remains intact however to regenerate and heal the burned area quickly, usually without scarring. Broken blisters provide possible infection sites prior to healing.

Third degree burns cause tissue death all the way through the skin, including the stem cells required to regenerate skin tissue. The only way a 3rd degree burn can heal is by skin regrowth from the edges, a slow process that usually results in scarring, unless skin grafts are used. Before healing 3rd degree burns present serious risk of infection, and can cause serious fluid loss. A 3rd degree burn over 25% of the body (or more) will typically precipitate shock in minutes, which itself requires prompt medical attention.

Even more serious burns are possible, which have been classified as fourth (even fifth) degree burns. These burns destroy tissue below the skin: muscle, connective tissue etc. They can be caused by thermal radiation exposures substantially in excess of those in the table for 3rd degree burns. Many people close to the hypocenter of the Hiroshima bomb suffered these types of burns. In the immediate vicinity of ground zero the thermal radiation exposure was 100 c/cm^2, some fifteen times the exposure required for 3rd degree burns, most of it within the first 0.3 seconds (which was the arrival time of the blast wave). This is sufficient to cause exposed flesh to flash into steam, flaying exposed body areas to the bone.

At the limit of the range for 3rd degree burns, the time lapse between suffering burns and being hit by the blast wave varies from a few seconds for low kiloton explosions to a minute of so for high megaton yields.

5.6.1.2 Incendiary Effects

Despite the extreme intensity of thermal radiation, and the extraordinary surface temperatures that occur, it has less incendiary effect than might be supposed. This is mostly due to its short

duration, and the shallow penetration of heat into affected materials. The extreme heating can cause pyrolysis (the charring of organic material, with the release of combustible gases), and momentary ignition, but it is rarely sufficient to cause self-sustained combustion. This occurs only with tinder-like, or dark, easily flammable materials: dry leaves, grass, old newspaper, thin dark flammable fabrics, tar paper, etc. The incendiary effect of the thermal pulse is also substantially affected by the later arrival of the blast wave, which usually blows out any flames that have already been kindled. Smoldering material can cause reignition later however.

The major incendiary effect of nuclear explosions is caused by the blast wave. Collapsed structures are much more vulnerable to fire than intact ones. The blast reduces many structures to piles of kindling, the many gaps opened in roofs and walls act as chimneys, gas lines are broken open, storage tanks for flammable materials are ruptured. The primary ignition sources appear to be flames and pilot lights in heating appliances (furnaces, water heaters, stoves, ovens, etc.). Smoldering material from the thermal pulse can be very effective at igniting leaking gas.

Although the ignition sources are probably widely scattered a number of factors promote their spread into mass fires. The complete suppression of fire fighting efforts is extremely important. Another is that the blast scatters combustible material across fire breaks that normally exist (streets, yards, fire lanes, etc.).

The effectiveness of building collapse, accompanied by the disruption of fire fighting, in creating mass fires can be seen in the San Francisco earthquake (1906), the Tokyo-Yokahama earthquake (1923), and the recent Kobe earthquake (1995). In these disasters there was no thermal radiation to ignite fires, and the scattering of combustible materials did not occur, but huge fires still resulted. In San Francisco and Tokyo-Yokohama these fires were responsible for most of the destruction that occurred.

In Hiroshima the fires developed into a true firestorm. This is an extremely intense fire that produces a rapidly rising column of hot air over the fire area, in turn powerful winds are generated which blow in to the fire area, fanning and feeding the flames. The fires continue until all combustible material is exhausted. Firestorms develop from multiple ignition sources spread over a wide area that create fires which coalesce into one large fire. Temperatures in firestorm areas can reach many hundreds of degrees, carbon monoxide reaches lethal levels, few people who see the interior of a firestorm live to tell about it. Firestorms can melt roads, cars, and glass. They can boil water in lakes and rivers, and cook people to death in buried bomb shelters. The in-blowing winds can reach gale force, but they also prevent the spread of the fires outside of the area in which the firestorm initially develops. The firestorm in Hiroshima began only about 20 minutes after the bombing.

Nagasaki did not have a firestorm, instead it had a type of mass fire called a conflagration. This is a less intense type of fire, it develops and burns more slowly. A conflagration can begin in multiple locations, or only one. Conflagrations can spread considerable distances from their origins. The fires at Nagasaki took about 2 hours to become well established, and lasted 4-5 hours.

5.6.1.3 Eye Injury

The brightness and thermal output of a nuclear explosion presents an obvious source of injury to the eye. Injury to the cornea through surface heating, and injury to the retina are both possible risks. Surprisingly, very few cases of injury were noted in Japan. A number of factors acted to reduce the risk. First, eye injury occurs when vision is directed towards the fireball. People spend relatively little time looking up at the sky so only a very small portion of the population would have their eyes directed at the fireball at the time of burst. Second, since the bomb exploded in bright daylight the eye pupil would be expected to be small.

About 4% of the population within the 3rd degree burn zone at Hiroshima reported keratitis, pain and inflammation of the cornea, which lasted several hours to several days. No other corneal damage was noted.

The most common eye injury was flashblindness, a temporary condition in which the visual pigment of retina is bleached out by the intense light. Vision is completely recovered as the pigment is regenerated, a process that takes several seconds to several minutes. This can cause serious problems though in carrying out emergency actions, like taking cover from the oncoming blast wave.

Retinal injury is the most far reaching injury effect of nuclear explosions, but it is relatively rare since the eye must be looking directly at the detonation. Retinal injury results from burns in the area of the retina where the fireball image is focused. The brightness per unit area of a fireball does not diminish with distance (except for the effects of haze), the apparent fireball size simply gets smaller. Retinal injury can thus occur at any distance at which the fireball is visible, though the affected area of the retina gets smaller as range increases. The risk of injury is greater at night since the pupil is dialated and admits more light. For explosions in the atmosphere of 100 kt and up, the blink reflex protects the retina from much of the light.

5.6.2 Blast Damage and Injury

Blast damage is caused by the arrival of the shock wave created by the nuclear explosion. Shock waves travel faster than sound, and cause a virtually instantaneous jump in pressure at the shock front. The air immediately behind the shock front is accelerated to high velocities and creates a powerful wind. The wind in turn, creates dynamic pressure against the side of objects facing the blast. The combination of the pressure jump (called the overpressure)and the dynamic pressure causes blast damage.

Both the overpressure and dynamic pressure jump immediately to their peak values when the shock wave arrives. They then decay over a period ranging from a few tenths of a second to several seconds, depending on the strength of the blast and the yield. Following the this there is a longer period of weaker negative pressure before the atmospheric conditions return to normal. The negative pressure has little significance as far as causing damage or injury is concerned. A given pressure is more destructive from a larger bomb, due its longer duration.

The is a definite relationship between the overpressure and the dynamic pressure. The overpressure and dynamic pressure are equal at 70 psi, and the wind speed is 1.5 times the speed of sound. Below an overpressure of 70 psi, the dynamic pressure is less than the overpressure; above 70 psi it exceeds the overpressure. Since the relationship is fixed it is convenient to use the overpressure alone as a yardstick for measuring blast effects. At 20 psi overpressure the wind speed is still 500 mph, higher than any tornado wind.

As a general guide, city areas are completely destroyed (with massive loss of life) by overpressures of 5 psi, with heavy damage extending out at least to the 3 psi contour. The dynamic pressure is much less than the overpressure at blast intensities relevant for urban damage, although at 5 psi the wind speed is still 162 mph - close to the peak wind speeds of the most intense hurricanes.

Humans are actually quite resistant to the direct effect of overpressure. Pressures of over 40 psi are required before lethal effects are noted. This pressure resistance makes it possible for unprotected submarine crews to escape from emergency escape locks at depths as great as one hundred feet (the record for successful escape is actually an astonishing 600 feet, representing a pressure of 300 psi). Loss of eardrums can occur, but this is not a life threatening injury.

The danger from overpressure comes from the collapse of buildings that are generally not as resistant. The violent implosion of windows and walls creates a hail of deadly missiles, and the collapse of the structure above can crush or suffocate those caught inside.

The dynamic pressure causes can cause injury by hurling large numbers of objects at high speed. Urban areas contain many objects that can become airborne, and the destruction of buildings generates many more. Serious injury or death can also occur from impact after being thrown through the air.

Blast effects are most dangerous in built-up areas due to the large amounts of projectiles created, and the presence of obstacles to be hurled against.

The blast also magnifies thermal radiation burn injuries by tearing away severely burned skin. This creates raw open wounds that readily become infected.

These many different effects make it difficult to provide a simple rule of thumb for assessing the magnitude of harm produced by different blast intensities. A general guide is given below:

1 psi Window glass shatters

Light injuries from fragments occur.

3 psi Residential structures collapse.

Serious injuries are common, fatalities may occur.

5 psi Most buildings collapse.

Injuries are universal, fatalities are widespread.

10 psi Reinforced concrete buildings are severely damaged or demolished.

Most people are killed.

20 psi Heavily built concrete buildings are severely damaged or demolished.

Fatalities approach 100%.

Suitable scaling constants for the equation r_blast = Y^0.33 * constant_bl

are:

constant_bl_1_psi = 2.2

constant_bl_3_psi = 1.0

constant_bl_5_psi = 0.71

constant_bl_10_psi = 0.45

constant_bl_20_psi = 0.28

where Y is in kilotons and range is in km.

5.6.3 Radiation Injury

Ionizing radiation produces injury primarily through damage to the chromosomes. Since genetic material makes up a very small portion of the mass of a cell, the damage rarely occurs from the direct impact of ionizing radiation on a genetic molecule. Instead the damage is caused by the radiation breaking up other molecules and forming chemically reactive free radicals or unstable compounds. These reactive chemical species then damage DNA and disrupt cellular chemistry in other ways - producing immediate effects on active metabolic and replication processes, and long-term effects by latent damage to the genetic structure.

Cells are capable of repairing a great deal of genetic damage, but the repairs take time and the repair machinery can be overwhelmed by rapid repeated injuries. If a cell attempts to divide before sufficient repair has occurred, the cell division will fail and both cells will die. As a consequence, the tissues that are most sensitive to radiation injury are ones that are undergoing rapid division. Another result is that the effects of radiation injury depend partly on the rate of exposure. Repair mechanisms can largely offset radiation exposures that occur over a period of time. Rapid exposure to a sufficiently large radiation dose can thus cause acute radiation sickness, while a longer exposure to the same dose might cause none.

By far the most sensitive are bone marrow and lymphatic tissues - the blood and immune system forming organs of the body. Red blood cells, which provide oxygen to the body, and white blood cells, which provide immunity to infection, only last a few weeks or months in the body and so must be continually replaced. The gastrointestinal system is also sensitive, since the lining of the digestive tract undergoes constant replacement. Although they are not critical for health, hair follicles also undergo continual cell division resulting in radiation sickness' most famous symptom - hair loss. The tissues least sensitive to radiation are those that never undergo cell division (i.e. the nervous system).

This also means that children and infants are more sensitive to injury than adults, and that fetuses are most sensitive of all.

If the individual survives, most chromosome damage is eventually repaired and the symptoms of radiation illness disappear. The repair is not perfect however. Latent defects can show up years or decades later in their effects on reproductive cells, and in the form of cancer. These latent injuries are a very serious concern and can shorten life by many years. They are the sole form of harm from low level radiation exposure.

5.6.3.1 Units of Measurement for Radiation Exposure

Three units of measurement have been commonly used for expressing radiation exposure: roentgens (R), rads, rems, the "three r's" of radiation measurement. In the scientific literature these are dropping out of use in favor of the SI (System Internationale) units grays (Gy) and sieverts (Sv). Each of the "three r's" measures something different. A rad is a measure of the amount of ionizing . A roentgen measures the amount of ionizing energy, in the form of energetic photons (gamma rays and x-rays) energy to which an organism is exposed. This unit is the oldest of the three and is defined more the convenience of radiation measurement, than for interpreting the effects of radiation on living organisms. Of more interest is the rad, since it includes all forms of ionizing radiation, and in addition measures the dose that is *actually absorbed* by the organism. A rad is defined as the absorption of 100 ergs per gram of tissue (or 0.01 J/kg). The gray measures absorbed doses as well, one gray equals 100 rads. The rem is also concerned with all absorbed ionizing radiations, and also takes into account the *relative effect* that different types of radiation produce. The measure of effect for a given radiation is its Radiation Biological Effect (RBE). A rem dose is calculated by multiplying the dose in rads for each type of radiation by the appropriate RBE, then adding them all up. The sievert is similar to the rem, but is derived from the gray instead of the rad. Sieverts use a somewhat simplified

system of measuring biological potency - the quality factor (Q). One sievert is roughly equal to 100 rems. The rem and the sievert are the most meaningful unit for measuring and discussing the effects of radiation injury.

Type Of Radiation	RBE	Q	
Gamma rays/X-rays	1	1	
Beta Particles	1	1	
Alpha Particles	10-20	20	(ingested emitter)
Neutrons (fast)	-	10	Overall effects

1	Immediate Effect
4-6	Delayed cataract formation
10	Cancer Effect
20	Leukemia Effect

5.6.3.2 Types of Radiation Exposure

An important concept to understand is the distinction between _whole body doses_ and radiation exposures concentrated in particular organs. The radiation dose units described above are defined per unit weight of tissue. An exposure of 1000 rems can thus refer to an exposure of this intensity for the whole body, or for only a small part of it. The total absorbed radiation energy will be much less if only a small part of the body is affected, and the overall injury will be reduced.

Not all tissues are exposed equally even in whole body exposures. The body provides significant shielding to internal organs, so tissues located in the center of the body may receive doses that are only 30-50% of the nominal total body dose rate. For example there is a 50% chance of permanent female sterility if ovaries are exposed to 200 rems, but this internal exposure is only encountered with whole body doses of 400-600 rems.

Radiation exposures from nuclear weapons occur on three time scales:

-
- The shortest is exposure from the prompt radiation emitted by the fireball which lasts about one minute. This can cause very intense exposures for individuals close to the

burst point. Neutron bombs rely on prompt radiation as the primary damage mechanism, in this case the prompt radiation arrives in a fraction of a second.

- The second scale is due to early (tropospheric) fallout from ground bursts. Fallout particles begin settling to the ground within an hour to a few hours after an explosion, most of the fallout descend within a day or two. At any particular site, the fallout deposition will last no more than several hours. Radiation exposure is accumulated as long as an individual remains within the fallout deposit zone, but due to the rapid initial decay most of the radiation exposure is incurred within the first few days. Exposures can be very large during the first few days.
- The third scale is long term exposure to low levels of radiation, lasting months or years. This may be due to any of several causes:
    - prolonged residence in areas contaminated by early fallout;
    - exposure to delayed (stratospheric) fallout;
    - exposure to radioisotopes absorbed by the body.
- Long term exposures are not intense, but large total doses can accumulate over long periods of time.

The effects of radiation exposure of usually divided into acute and latent effects. Acute effects typically result from rapid exposures, the effects show up within hours to weeks after a sufficient dose is absorbed. Latent effects take years to appear, even after exposure is complete.

Since the latent effects of radiation exposure are cumulative, and there does not appear to be any threshold exposure below which no risk is incurred, radiation safety standards have been set to minimize radiation exposure over time. Current standards are:

Occupational Exposure

0.3 rem/wk (whole body exposure)

1.5 rem/yr (whole body exposure for pregnant women)

5 rem/yr (whole body exposure)

15 rem/yr (eye tissue exposure)

50 rem/yr (limit for any tissue)

200 rem lifetime limit (whole body exposure)

Public Exposure

0.5 rem/yr (whole body exposure)

5 rem/yr (limit for any tissue)

The occupational exposure limits are likely to be reduced soon (if they have not been already).

The normal human annual radiation exposure varies considerably with location (elevation and surface mineral composition), and medical treatment. Typical values are 0.1 rems from natural radiation and 0.08 rems from medical x-rays, for a total of 0.18 rem/yr. In the US, Colorado has one of the highest natural backgrounds (0.25 rem) since high altitudes cause greater cosmic ray exposures, and granite rock formations contain uranium series radioisotopes. If natural radioisotopes are unusually concentrated, levels as high as 0.5-12 rems/yr have been recorded (some areas of Sri Lanka, Kerala India, and Brazil). This does not count indoor radon exposure which depends heavily on building design, but can easily exceed all other exposure sources combined in regions with high soil radon levels. This source has been known to cause lung exposures in the home of 100 rem/yr (a risk factor comparable to heavy smoking)!

5.6.3.3 Prompt Radiation Emission From Nuclear Explosions

Although the subject is complex, a simplified guide to estimating the prompt radiation exposure from nuclear explosions is given here. The following scaling law can be used to determine the lethal radius with yield:

r_radiation = Y^0.19 * constant_rad

If Y is in kilotons, range is in meters, and the dose standard is 1000 rads then:

constant_rad_1000 = 700 m

This can then be scaled for distance by adjusting for attenuation with range using the table below. The table lists tenth-ranges, the distance over which the dose decreases (for greater distance) or increases (for shorter distance) by a factor of 10.

1 kt 330 m

10 kt 440 m

100 kt 490 m

1 Mt 560 m

10 Mt 670 m

20 Mt 700 m

So, for example to calculate the radiation dose for a 10 Mt bomb at 5000 m, we calculate: dose = (1000 rads) / 10^[(5000-[10000^0.19]*700)/670] = 35 rads

This guide assumes 100% fission yield for bombs <100 kt, and 50/50 fission/fusion for higher yields. Due to the enhanced radiation output of low-yield neutron bombs different factors need to be used:

constant_rad_1000 = 620 m

tenth-range 385 m

5.6.3.4 Acute Radiation Sickness

This results from exposure to a large radiation dose to the whole body within a short period of time (no more than a few weeks). There is no sharp cutoff to distinguish acute exposures from chronic (extended) ones. In general, higher total doses are required to produce a given level of acute sickness for longer exposure times. Exposures received over a few days do not differ substantially from instantaneous ones, except that the onset of symptoms is correspondingly delayed or stretched out. Nuclear weapons can cause acute radiation sickness either from prompt exposure at the time of detonation, or from the intense radiation emitted by early fallout in the first few days afterward.

The effects of increasing exposures are described below. A notable characteristic of increasing doses is the non-linear nature of the effects. That is to say, a threshold exists below which observable effects are slight and reversible (about 300 rems), but as exposures rise above this level the possibility of mortality (death) begins and increases rapidly with dose. This is believed to be due in part to the saturation of cellular repair mechanisms.

The total energy absorbed by a 75 kg individual with a whole body exposure of 600 rads (fatal in most cases) is 450 joules. It is interesting to compare this to the kinetic energy of a .45 caliber bullet, which is about 900 joules.

A power law for scaling radiation effects for longer term exposures has been proposed in which the dose required for a given effect increases by t^0.26, where time is in weeks. For exposures of one week or less the effect of rem of radiation is assumed to be constant. Thus an exposure capable of causing 50% mortality is 450 rems if absorbed in a week or less, but is 1260 rems if it occurs over a year.

5.6.3.4.1 Acute Whole Body Exposure Effects

Below 100 REMS

In this dose range no obvious sickness occurs. Detectable changes in blood cells begin to occur at 25 rems, but occur consistently only above 50 rems. These changes involve fluctuations in the overall white blood cell count (with drops in lymphocytes), drops in platelet counts, and less severe drops in red blood cell counts. These changes set in over a period of days and may require months to disappear. They are detectable only by lab tests. At 50 rems atrophy of lymph glands becomes noticeable. Impairment to the immune system could increase the susceptibility to disease. Depression of sperm production becomes noticeable at 20 rems, an exposure of 80 rems has a 50% chance of causing temporary sterility in males.

100-200 REMS

Mild acute symptoms occur in this range. Tissues primarily affected are the hematopoietic (blood forming) tissues, sperm forming tissues are also vulnerable. Symptoms begin to appear at 100 rems, and become common at 200 rems. Typical effects are mild to moderate nausea (50% probability at 200 rems) , with occasional vomiting, setting in within 3-6 hours after exposure, and lasting several hours to a day. This is followed by a latent period during which symptoms disappear. Blood changes set in and increase steadily during the latency period as blood cells die naturally and are not replaced. Mild clinical symptoms return in 10-14 days. These symptoms include loss of appetite (50% probability at 150 rems), malaise, and fatigue (50% probability at 200 rems), and last up to 4 weeks. Recovery from other injuries is impaired and there is enhanced risk of infection. Temporary male sterility is universal. The higher the dosage in this range, the more likely the effects, the faster symptoms appear, the shorter the latency period, and the longer the duration of illness.

200-400 REMS

Illness becomes increasingly severe, and significant mortality sets in. Hematopoietic tissues are still the major affected organ system. Nausea becomes universal (100% at 300 rems), the incidence of vomiting reaches 50% at 280 rems. The onset of initial symptoms occurs within 1-6 hours, and last 1-2 days. After this a 7-14 day latency period sets in. When symptoms recur, the may include epilation (hair loss, 50% probability at 300 rems), malaise, fatigue, diarrhea (50% prob. at 350 rems), and hemorrhage (uncontrolled bleeding) of the mouth, subcutaneous tissue and kidney (50% prob. at 400 rems). Suppression of white blood cells is severe, susceptibility to infection becomes serious. At 300 rems the mortality rate without medical treatment becomes substantial (about 10%). The possibility of permanent sterility in females begins to appear. Recovery takes 1 to several months.

400-600 REMS

Mortality rises steeply in this dose range, from around 50% at 450 rems to 90% at 600 (unless heroic medical intervention takes place). Hematopoietic tissues remain the major affected organ system. Initial symptoms appear in 0.5-2 hours, and last up to 2 days. The latency period remains 7-14 days. The symptoms listed for 200-400 rems increase in prevalence and severity,

reaching 100% occurrence at 600 rems. When death occurs, it is usually 2-12 weeks after exposure and results from infection and hemorrhage. Recovery takes several months to a year, blood cell counts may take even longer to return to normal. Female sterility becomes probable.

600-1000 REMS

Survival depends on stringent medical intervention. Bone marrow is nearly or completely destroyed, requiring marrow transfusions. Gastrointestinal tissues are increasingly affected. Onset of initial symptoms is 15-30 minutes, last a day or two, and are followed by a latency period of 5-10 days. The final phase lasts 1 to 4 weeks, ending in death from infection and internal bleeding. Recovery, if it occurs, takes years and may never be complete.

Above 1000 REMS

Very high exposures can sufficient metabolic disruption to cause immediate symptoms. Above 1000 rems rapid cell death in the gastrointestinal system causes severe diarrhea, intestinal bleeding, and loss of fluids, and disturbance of electrolyte balance. These effects can cause death within hours of onset from circulatory collapse. Immediate nausea occurs due to direct activation of the chemoreceptive nausea center in the brain.

In the range 1000-5000 rems the onset time drops from 30 minutes to 5 minutes. Following an initial bout of severe nausea and weakness, a period of apparent well-being lasting a few hours to a few days may follow (called the "walking ghost" phase). This is followed by the terminal phase which lasts 2-10 days. In rapid succession prostration, diarrhea, anorexia, and fever follow. Death is certain, often preceded by delirium and coma. Therapy is only to relieve suffering.

Above 5000 rems metabolic disruption is severe enough to interfere with the nervous system. Immediate disorientation and coma will result, onset is within seconds to minutes. Convulsions occur which may be controlled with sedation. Victim may linger for up to 48 hours before dying.

The U.S. military assumes that 8000 rads of fast neutron radiation (from a neutron bomb) will immediately and permanently incapacitate a soldier.

It should be noted that people exposed to radiation doses in the 400-1000 rem range following the Chernobyl disaster had much higher rates of survival than indicated above. This was made possible by advances in bone marrow transfusions and intensive medical care, provided in part by Dr. Robert Gale. However two caveats apply:

Such care is only available if the number of cases is relatively small, and the infrastructure for providing it is not disrupted. In the case of even a limited nuclear attack it would be impossible to provide more than basic first aid to most people and the fatality rates might actually be higher than given here.

Many of the highly exposed Chernobyl survivors have since died from latent radiation effects.

5.6.3.4.2 Acute Localized Tissue Exposure

Localized acute exposure is important for two organs: the skin, and the thyroid gland.

Beta Burns

Beta particles have a limited range in tissue. Depending on their energy, betas are completely absorbed by 1 mm to 1 cm of tissue. External exposures to beta particles from fallout thus primarily affect the skin, causing "beta burns". Due to the poor penetrating power of betas, these injuries only occur if there is direct skin exposure to fallout particles, or if an individual remains outdoors in a strong radiation field. Remaining indoors, wearing substantial clothing, and decontamination by washing can prevent this type of exposure. Beta burns were encountered in Marshall Islanders, and the crew of a Japanese fishing vessel, following the Castle Bravo test which unexpectedly dumped high fallout levels over a large area.

The initial symptom for beta burns are an itching or burning sensation during the first 24-48 hours. These symptoms are marked only if the exposure is intense, and do not occur reliably. Within 1-2 days all symptoms disappear, but after 2-3 weeks the burn symptoms appear. The first evidence is increased pigmentation, or possibly erythema (reddening). Epilation and skin lesions follow.

In mild to moderate cases damage is largely confined to the epidermis (outer skin layers). After forming a dry scab, the superficial lesions heal rapidly leaving a central depigmented area, surrounded by an irregular zone of increased pigmentation. Normal pigmentation returns over a few weeks.

In more serious cases deeper ulcerated lesions form. These lesions ooze before becoming covered with a hard dry scab. Healing occurs with routine first aid care. Normal pigmentation may take months to return.

Hair regrowth begins 9 weeks after exposure and is complete in 6 months.

Thyroid Exposure

The short-lived radioisotope iodine-131 (half-life 8 days) presents a special risk due to the tendency for ingested iodine to be concentrated in the thyroid gland. This risk is mitigated by the fact that direct ingestion of fallout is rare, and easily avoided. Iodine-131 typically enters the body through the consumption of contaminated milk, which in turn results from milk cows consuming contaminated fodder.

The short half-life means that the initial radiation intensity of I-131 is high, but it disappears quickly. If uncontaminated fodder can be provided for a month or two, or if dry or canned milk can be consumed for the same period, there is little risk of exposure.

If I-131 contaminated food is consumed, about one-third of the ingested iodine is deposited in the thyroid gland which weighs some 20 g in adults, and 2 g in infants. This can result in very high dose rates to the gland, with negligible exposures to the rest of the body. Due to the smaller glands of infants and children, and their high dairy consumption, they are particularly vulnerable to thyroid injury. Some Marshallese children received thyroid doses as high as 1150 rems. Most of the children receiving doses over 500 rems developed thyroid abnormalities within 10 years, including hypothyroidism and malignancies.

I-131 exposure can be prevented by prompt consumption of potassium iodide supplements. Large doses of potassium iodide saturate the body with iodine and prevent any subsequent retention of radioiodine that is consumed. *5.6.3.4.3 Fetal Injury*

Acute radiation exposure during pregnancy can cause significant harm to the fetus. At Hiroshima and Nagasaki adverse effects were seen when pregnant women who were exposed

to 200 rems of radiation or more. When exposure occurred during the first trimester a significant increase in mentally impaired children were noted. When exposure occurred during the last trimester, there was a marked increase in stillbirths and in elevated infant mortality during the first year of life.

5.6.3.5 Chronic Radiation Exposure

Long term radiation exposure results from residing in a fallout contaminated area for an extended period (external exposure), consuming food produced in a contaminated area (internal exposure), or both. If the exposure rate is low enough, no symptoms of radiation sickness will appear even though a very large total radiation dose may be absorbed over time. Latent radiation effects (i.e. cancer, genetic damage) depend on total dosage, not dose rate, so serious effects can result. An exposure of 0.25 rem/day over 5 years would accumulate 450 rems with little chance of overt sickness, but it would have a high mortality rate if the exposure were acute. The exposure time scaling law given above also indicates that a slow onset of symptoms characteristic of acute radiation sickness can occur. As an example, the most heavily contaminated location of the Rongelap atoll (160 km downwind of the March 1, 1954 15 Mt Castle Bravo test), received a total accumulated exposure of 3300 rads. Of this, 1100 rads was accumulated during the interval from 1 month to 1 year following the test. If the site had been occupied during this period, the effective exposure for radiation sickness effects would be 1100/(48 weeks)^0.26 = 403 rads.

5.6.3.5.1 External Exposure

When an area is contaminated by gamma emitting isotopes, a radiation field is created that exposes all organisms that are not shielded from it. Only gamma rays have the necessary range and penetration to create a significant hazard. The principal source of long-term external exposure is cesium-137 (30 year half-life, 0.6 MeV gamma energy).

A megaton of fission yield produces enough Cs-137 to contaminate 100 km^2 with a radiation field of 200 rad/year. A megaton-range ground burst can contaminate an area of thousands of square kilometers with concentrations that would exceed occupational safety guidelines. 3,000 megatons of fission yield, if distributed globally by stratospheric fallout, would double the world's background radiation level from external exposure to this isotope alone.

It is possible to substantially reduce external exposure in contaminated areas by remaining indoors as much as possible. Exposure can be reduced by a factor of 2-3 for a frame house, or 10-100 for a multi-story building, and adding additional shielding to areas where much time is spent (like the bedroom) can increase these factors substantially. Since the half-life of Cs-137 is long, these would be permanent lifestyle adjustments. Such measures have been necessary (especially for children) in areas of Belarus that were heavily contaminated by Chernobyl.

5.6.3.5.2 Internal Exposure

Internal exposure to radiation is the most serious chronic risk from fallout if food grown in contaminated areas is consumed. Widespread contamination from a nuclear war, or a major radiation accident (like the Kyshtym and Chernobyl disasters), may leave no other practical choice. Alternatively, people residing in contaminated areas may come to disregard safety instructions about locally produced food (as has happened in the Marshall Islands and Ukraine). Radioisotopes may be taken up into plants through the root system, or they may be contaminated by fallout descending on the leaves. Gross contamination of food plants or fodder from the fallout plume of a ground burst is an obvious hazard, but the gradual descent of

worldwide fallout is also a problem.
The primary risks for internal exposure are cesium-137 and strontium-90. Strontium-89, transuranics alpha emitters, and carbon-14 are also significant sources of concern.
Only a few curies of radioisotopes per km^2 are sufficient to render land unsuitable for cultivation under current radiation safety standards. A megaton of fission yield can thus make some 200,000 km^2 useless for food production for decades. Depression of leukocyte levels have been observed in people in Belarus living in areas that were contaminated with only 0.2 curies/km^2.

Cesium-137

This alkali metal has chemistry resembling that of potassium. As a result, it is readily absorbed by food plants, and by animal tissues. Once consumed cesium distributes itself fairly evenly through the body, which means that Cs-137 absorption causes whole body exposure (a fact further aided by the penetrating nature of its gamma emissions). Cesium has a moderate residence time in the body, the residence half-life ranging from 50-100 days, so that the body will be cleared of the isotope once consumption of contaminated material ceases in a matter of several months, to a few years.

Strontium 90 and 89

Strontium is chemically similar to calcium, and is deposited in bone along with calcium. Most of the strontium ingested does not end up in bone, it has a biological half-life only 40 days. Somewhat less than 10% of the Sr is retained in the bone, but it has a biological half-life of 50 years. Since the bone marrow is among the most sensitive tissue in the body to radiation, this creates a very serious hazard.
Sr-90 (28.1 yr half-life) thus can cause long term damage, while Sr-89 (52 days) can cause significant short term injury. Safety exposure standards impose a Sr-90 body burden limit of 2 microcuries (14 nanograms) for occupational exposure, 0.2 microcuries for individual members of the general population, and 0.067 microCi averaged over the whole population. It is estimated that 10 microCi per person would cause a substantial rise in the incidence of bone cancer. The explosion of several thousands of fission megatons in the atmosphere could raise the average body burden of the entire human race to above the occupational exposure limit for Sr-90 for a couple of generations. Contamination of 2 curies of Sr-90 per km^2 is the U.S. limit for food cultivation.
Alpha emitting heavy elements can be serious health risks also. The isotopes of primary concern here are those present in substantial quantities in nuclear weapons: short lived uranium isotopes (U-232 and U-233) and transuranic elements (primarily Pu-239, Pu-240, and Americium-241). These elements are hazardous if ingested due to radiotoxicity from the highly damaging alpha particles. The quantities of these isotopes present after a nuclear explosion are negligible compared to the amount of fission product radioisotopes. They represent a hazard when nuclear weapons are involved in "broken arrow" incidents, that is, accidents where the fissile isotopes inside are released. The exposure areas are of course small, compared to the areas threatened by fallout from a nuclear detonation. A typical nuclear weapon will contain some 300-600 curies of alpha emitter (assuming 5 kg plutonium). The isotope breakdown is approximately: 300 curies Pu-239, 60 curies Pu-240, and up to 250 curies of Am-241.
If small particles of alpha emitters are inhaled, they can take up permanent residence in the lung and form a serious source of radiation exposure to the lung tissue. A microcurie of alpha emitter

deposited in the lungs produce an exposure of 3700 rems/yr to lung tissue, an extremely serious cancer risk.

Uranium and the transuranic elements are all bone-seekers (with the exception of neptunium). If absorbed, they are deposited in the bone and present a serious exposure risk to bone tissue and marrow. Plutonium has a biological half-life of 80-100 years when deposited in bone, it is also concentrated in the liver with a biological half-life of 40 years. The maximum permissible occupational body burden for plutonium-239 is 0.6 micrograms (.0375 microcuries) and 0.26 micrograms for lung burden (0.016 microCi).

Carbon-14 is a weak beta particle emitter, with a low level of activity due to its long half-life. It presents a unique hazard however since, unlike other isotopes, it is incorporated directly into genetic material as a permanent part throughout the body. This means that it presents a hazard out of proportion to the received radiation dose as normally calculated. *5.6.3.5.3 Cancer*

The most serious long term consequence of radiation exposure is the elevation of cancer risk. Estimates of the carcinogenicity of radiation, especially of low exposures, have tended to increase over the years as epidemiological data has accumulated.

The current state-of-the-art in low level risk estimation is the 1990 report issued by the National Academy of Sciences Committee on Biological Effects of Ionizing Radiation (BEIR) entitled _Health Effects of Exposure to Low Levels of Ionizing Radiation_, also known as BEIR V.

As a general rule of thumb, it appears that cancer risk is more or less proportional to total radiation exposure, regardless of the quantity, rate or duration. 500 rems received over a decade is thus as serious a risk as 500 rems received all at once, and 50 rems is one-tenth as bad as 500. There is no evidence of a threshold effect or "safe dose". Safety standards are established primarily to keep the increased incidence of cancer below detectable levels.

Significant deviations from the above rule of proportionality for total exposure do occur. In particular, low doses (for which the risk is small anyway) received over an extended period of time are significantly less carcinogenic (by about a factor of 2) than the same dose received all at once.

Cancer risk to radiation exposure can be expressed as the increase in the lifetime probability of contracting fatal cancer per unit of radiation. The current estimate of overall risk is about a 0.8% chance of cancer per 10 rems for both men and women, averaged over the age distribution of the U.S. population. Thus a 1000 rem lifetime whole body radiation exposure would bring about a 80% chance of contracting fatal cancer, in addition to the normal incidence of cancer (about 20%). The risk for children appears to be about twice as great (due at least partly to the fact that they will live longer after exposure, and thus have greater opportunity to contract cancer).

There are also risk coefficient for specific tissue exposures. These are (approximately):

Female Breast 1.0%/100 rems

Bone Marrow 0.2%/100 rems (0.4% for children)

Bone Tissue 0.05%/100 rems

Lung 0.2%/100 rems

5.6.3.5.4 Genetic Effects

Radiation damage to the germ cells of the reproductive organs can cause mutations that are passed on to subsequent generations. Although this is very important, it can nonetheless be

overplayed. It may seem surprising, but no elevated mutation rate from radiation has ever been detected in the human population, not even in the substantial population of atomic bomb survivors and descendants. One reason for this is that humans are wild animals, that is, they have not been subjected to controlled breeding and thus have a high incidence of natural genetic variability and disorders, compared to laboratory and domestic animals. About 10% of the human population has detectable genetic disorders (most are not serious). This makes it difficult to detect additional mutations unless the rate is also high.
Two factors act to limit the effective radiation exposure for genetic effects, one for acute exposures, the other for chronic exposures. High acute exposures to the reproductive organs can cause permanent sterility, which prevents transmission of genetic effects. The cumulative effect of chronic exposure is limited by the fact that only exposures prior to reproduction count. Since most reproduction occurs before the age of 30, exposures after that age have little effect on the population.
It is estimated that the dose to reproductive tissue required to double the natural incidence of genetic disorders is 100-200 rems. The initial rate of observable disorders (the first generation) is only about 1/3 of the eventual rate once genetic equilibrium is established. Of course increases in the rate of genetic disorders (especially in a large population) is a _permanent_ alteration of the human species.
5.6.3.5.5 Cataracts
Eye tissue exposed to radiation shows an increased incidence of cataracts at dose levels below which most tissues show increased cancer rates. This makes cataract risk the most important tissue dose criterion for establishing safety standards.

6.0 Nuclear Materials

This section describes the basic facts and properties about the materials that are uniquely important to applications involving the large scale release of nuclear power.

- 6.1 Production of Isotopes
- 6.2 Fissionable Materials
- 6.3 Fusionable Materials
- 6.4 Other Materials

6.1 Production of Isotopes

The most critical materials required for nuclear weapons are special isotopes of particular elements. Some of these isotopes exist in nature, but are highly diluted by other isotopes of the same element (e.g. deuterium, Li-6, U-235). These isotopes must be enriched (concentrated) to be useful. Other isotopes naturally exist only in minuscule quantities, if at all, and must be manufactured through nuclear reactions (e.g. tritium and Pu-239). An overview of these two classes of techniques are given below.

6.1.1 Isotopic Enrichment

All isotopes of the same element have nearly identical chemical properties, consequently separating these isotopes must rely on subtle effects created by the differences in mass of the atoms. Some of these techniques are physical - they rely on physical processes that have mass dependent effects. Others are chemical - isotopic mass differences can affect the rate of chemical reactions.

Since small effects are being used to achieve separation, the achievable concentration increase for any enrichment process tends to be small. This can be overcome by applying the enrichment process many, many times in a series of enrichment steps or "stages". Each step takes the enriched product of the previous step and enriches it further. This creates a sequential enriching system called a cascade.

There are two important factors that affect the performance of a cascade. First is the separation factor, which is a number greater than one, that measures the inherent enrichment capability of a single stage. The lower the separation factor, the more stages required to achieve the desired degree of concentration. The second factor is the allowable loss of the desired isotope in the waste stream.

To understand this second factor, note that at each stage the material flow feeding the process is split into two parts: the enriched product, and the depleted waste. Since the enrichment in the product is small, so is the depletion in the waste. The amount of desirable isotope in the waste stream can easily exceed the amount in the enriched product. If the isotope being concentrated is rare, this is very undesirable. Also if the feed stock is already significantly enriched, then the waste stream is also enriched and too valuable to simply throw away.

These considerations demand that the cascade be constructed so that the waste from downstream (later) enrichment stages be fed back to earlier stages. Each stage then receives feed stock consisting of the enriched product from the next upstream stage, and the waste from the next downstream cascade.

The starting material usually does not enter at the beginning of the cascade (that is, the stage that is farthest upstream), it enters farther down. The stages that are upstream from the starting feed (called the tail of the cascade) strip the waste stream of the desired isotope so that the waste that actually leaves the cascade is greatly depleted. The stages making up the head of the cascade produce progressively more enriched product until the most highly enriched product exists the last stage.

For any particular cascade the actual amount of enriched product produced depends on the degree of enrichment desired AND on the degree of depletion in the waste stream. If it is okay to reject as waste most of desired isotope that enters the cascade, then the output of the cascade can be increased by producing a large waste stream of slightly depleted material (this can be done if the starting material is cheap). If the starting material is precious, then a high percentage of the desired isotope can be stripped - which produces highly depleted waste.

Examples of these two extremes are deuterium enrichment where the starting material is water and the feedstock cost is essentially zero, and uranium enrichment where the feedstock cost is substantial.

There is quite an assortment of isotope enrichment processes that have been used (or proposed) to produce substantial amounts of enriched isotopes. A sampling of important ones are:

- Electromagnetic separation
- Gaseous diffusion
- Thermal liquid diffusion
- Gas Centrifuge Separation
- Aerodynamic separation
- AVLIS (atomic vapor laser isotope separation)
- Chemical exchange
- Distillation
- Electrolysis

The effectiveness of any technique varies from element to element. Some of the techniques listed above can be applied to almost any element (e.g. electromagnetic separation), others are useful for only a small number of isotopes (e.g. electrolysis).

The separation factor for many of these processes (diffusion, chemical exchange, distillation) depends on the mass ratio between the isotopes. These techniques work best for light elements since each unit of atomic weight difference has a bigger effect. The largest isotope mass ratio (for naturally occurring isotopes) is 2, for deuterium and hydrogen-1.

Other processes depend on the absolute mass difference, or are otherwise insensitive to mass ratios (electromagnetic, centrifuge, aerodynamic, and AVLIS separation techniques). These work well for heavy elements, and are the methods of choice for uranium enrichment.

The cost of an enrichment process is mostly dependent on two factors: the capital cost of the plant, and the energy consumption for operation. Low separation factors require more cascade stages and thus both larger capital and operating costs. The capital cost is roughly proportional to 1/(s - 1)^2 where s is the separation factor. Different separation techniques differ dramatically in the amount of energy required for a single stage, so a low energy cost technique can offset a larger number of stages, or a higher cost per stage.

Another factor of some interest is the cascade equilibrium time. Due to limitations in the output rate of each cascade unit, the amount of material that is passed back to earlier stages by each unit, and the fact that each unit holds a certain amount of material, it takes a significant amount of time when a cascade is started for the final product to emerge and reach a steady output rate. Some types of cascades can take a year to reach equilibrium.

6.1.1.1 Electromagnetic Separation

This technique uses magnetic fields to separate a stream of ionized atoms in a device called a calutron. This is basically a scaled up mass spectrometer. The different masses of the isotopes give each a different radius of curvature in the magnetic field, causing the stream to separate into two (or more). The streams are then captured by target cells. Electromagnetic separation is the most versatile technique and can be used with any combination of isotopes. It also has an extremely large separation factor. Typically two passes are sufficient to produce high purity isotopes (>80%) from low concentration starting material (less than 1%), and will also efficiently strip the desired isotope from the waste stream. However it is an inefficient production process, since most of the feed is lost inside the calutron and must be recovered during periodic maintenance work. Other shortcomings of electromagnetic separation are the enormous energy consumption, the difficult and costly equipment maintenance, and the limited production rate given the capital cost. It is mainly used for producing small quantities of pure isotopes for research.

6.1.1.2 Gaseous Diffusion

This technique is practical for isotopes that have a substantial mass difference, and can be converted to gaseous form (and important caveat, the mass difference of the gas molecules should be due only to the mass difference of the isotopes to be separated). The technique works by allowing the gas to diffuse through a porous barrier. Lighter, faster molecules diffuse slightly quicker. The quality of the barrier material is critical. It must have even pore size, have no leaks at all, and be able to withstand substantial pressure differences, and chemical attack by the gaseous isotope. Aside from the barrier material problem, gaseous diffusion requires rather ordinary chemical plant technology (by today's standards). For some light elements the separation factor can be fairly large, but for most it is small (1.00429 for uranium). A large number of cascade stages can still achieve high purity, although the plant size must be large, and the cost of construction and operation is high. This technique has been the dominant enrichment process for uranium since the first plant was built in 1945. Although gas diffusion has been displaced by gas centrifuges for new plants since the late 70s, the large previously installed base makes this still the largest enrichment process worldwide.

6.1.1.3 Thermal Liquid Diffusion

This method makes use of the fact that different molecular velocities cause lighter molecules to concentrate in hotter regions when a temperature gradient exists. The separation factor is determined by the ratio of the mass difference and the total masses, so it is larger for light elements. In general the separation factor is small (around 1.01), so many cascade stages are needed. Although the equipment required is simple, the power requirement (for heat) is extremely large, although this can be reduced if waste heat is available from other industrial processes. This prevents it from being practical on a large scale.

6.1.1.4 Gas Centrifuge Separation

Although Germany developed this technology during W.W.II, it did not see application until the fifties and sixties. This technique works by passing gaseous isotope mixture through high speed

centrifuges. Centrifugal force causes the isotopes to concentrate in lighter and heavier gas layers which are then bled off. The major advantage of gas centrifuges is that the separation factor is due to absolute mass differences, not to mass ratios. It works as well with heavy elements as light ones. Separation factors are proportional to the square of the ratio between the rotational speed and the molecular velocity in the gas. The highest speeds possible are thus very desirable. Rotational speeds above the speed of sound (>330 m/sec) are typical, advanced designs using carbon fiber rotors can exceed 600 m/sec (in principal Kevlar rotors could reach speeds as high as 1100 m/sec). Typical separation factors can range from 1.01 to over 1.10. Separation factors depend on the absolute mass difference between isotopes (not the mass ratio) and the square of the peripheral speed. Thus compared with gaseous diffusion, each centrifuge produces a much larger separation factor for heavy elements. A series of separation stages are still needed, but there can be many fewer. Gas centrifuges consume much less energy than gaseous diffusion, and it is much easier to adjust the scale of production plants. Gas centrifuges are the dominant uranium enrichment process for new plants in the world today. A high level of technical proficiency is required to build efficient centrifuges, which delayed their commercial introduction (along with competition from government-subsidized gaseous diffusion plants).

6.1.1.5 Aerodynamic Separation

This can be thought of as a variant on the gas centrifuge approach, except that instead of spinning a chamber through which a gas is passed, the gas is spun inside a nozzle or vortex tube by injecting it under high pressure. It is considerably more energy intensive than centrifuge separation and appears at least as intensive as gaseous diffusion. This technology has been developed by South Africa and Germany for uranium enrichment.

6.1.1.6 AVLIS (atomic vapor laser isotope separation)

This technology makes use of the fact that isotopes of different masses absorb slightly different wavelengths of light (an indirect consequence of the nucleus mass difference). Precisely tuned lasers would excite only the isotope atoms desired in a stream of atomic vapor. The ionized atoms would then be separated from the neutral ones electromagnetically or by chemical reaction. AVLIS has not been used on an industrial scale yet. This technique promises to allow high efficiency production of high-purity U-235 and Pu-239, although its true useful is difficult to judge without an operating plant to observe. AVLIS technology, if available, could make it possible for a country to produce substantial batches of weapon-grade uranium or plutonium (from commercial reactor waste) without being detected. The energy required for separation itself is very low, only enough to break the molecular bond or ionize the atom. Energy consumption is mostly determined by the efficiency of the laser used, which is generally on the order of 0.1%.

6.1.1.7 Chemical Exchange

This technique makes use of different chemical reaction rates between isotopes. It works best for light elements where the reaction rate differences are large. Practical plants use reactions

that allow the two reactants to be in different phases (gas/liquid, solid/liquid, immiscible liquids). This permits convenient separation of the enriched and depleted materials, and allows continuous countercurrent operation. By also including temperature differences between the phases, the separation factor can be considerably enhanced. This is the most important process today for producing heavy water, for which it is by far the most energy-efficient method. Chemical exchange techniques for uranium have also been developed by Japan and France, but have never been used for production work. Chemical exchange is also used for lithium-6 enrichment.

6.1.1.8 Distillation

This is another thermal concentration technique, light isotopes evaporate more rapidly than heavier ones. Again, it works best for light elements with large mass ratios. It also helps if the heat of vaporization of the material being distilled is low. This has been applied successfully to heavy water production.

6.1.1.9 Electrolysis

This technique is unique to the production of heavy water. For reasons that are not entirely understood, when water is electrolyzed deuterium is concentrated with the remarkable separation factor of 7 or even more. However, electrolysis is so energy intensive it is only practical for processing enriched water in the later stages of heavy water production.

Representative Separation Factors

	H/D	C-12/13	U-235/238
Chemical Exchange	1.2-3	1.02	1.0015
Distillation	1.05-1.6	1.01	nil
Gaseous Diffusion	1.2	1.03	1.00429
Centrifuge (250 m/sec)	1.01	1.01	1.026
Centrifuge (600 m/sec)			1.233
Electrolysis	7		

The total number of stages in a cascade is given by:

Ln[R[N_p]/R[N_w]]/Ln[s] - 1

where N_p and N_w are the isotope concentrations in the feed and waste,

s is the separation factor, and

R[N] = N/(1 - N).

6.1.2 Transmutation

The transmutation of one isotope or element into another requires an extraordinary type of reagent: sub-atomic particles. Large nuclear weapon programs consume some of these artificial isotopes in ton quantities. Only one sub-atomic particle is available in the necessary kilogram amounts: neutrons. A 1 gigawatt nuclear reactor will produce about 3.75 kg of neutrons during the course of a year (assuming constant operation). The fact that neutrons have no electrical charge is very important since it allows them to easily penetrate the charged electron shells and nucleus of an atom.

Nuclear reactors produce copious quantities of neutrons by their very nature. Each fission produces roughly 2.5 neutrons, one of these is required to keep the reaction going, the others are available for transmuting other materials although some portion of these are unavoidably captured by materials making up the reactor structure. Reactors were the first source, and so far remain the only source, of large amounts of neutrons. They are not the only conceivable source though. At least two others have been proposed.

The first of these was simply to use a nuclear bomb to provide large amounts of neutrons. By incorporating the starting material into the bomb structure, or packing it around the bomb, substantial quantities of material could be transmuted in an instant. The shortcomings of this scheme are collecting the product, and avoiding undesirable side effects. Exploding the bomb underground could overcome these problems, the product would be removed from the rubble filled cavern left behind by mining. The difficulties with this process are formidable however, and nuclear explosions are themselves very unpopular no matter how well contained so this approach will probably never be used.

An alternative seriously proposed recently is to use particle accelerators to produce neutrons without reactors. Recent advances in accelerator technology make it possible to produce very intense beams of protons relatively efficiently. The protons, upon striking a target material of moderate atomic number with energies of a few hundred MeV, would knock loose some 50 neutrons each. Considerable time and money would be required to make this practical of course.

The principal isotopes produced by neutron transmutation are plutonium (produced from U-238) and tritium (from lithium-6). U-233 from thorium is also of considerable interest. Plutonium is produced as a byproduct by most nuclear reactors. Practical tritium production on the other hand requires special isotope production reactors. This is because most reactors use low-enriched uranium as fuel, which is mostly U-238 (96-97%). U-238 is thus normally present in the reactor in large amounts, and it absorbs most of the excess neutrons leaving few available for other purposes. Any natural or low-enriched uranium reactor could be considered a plutonium breeder, although the quality of the plutonium produced varies considerably depending on how the fuel and reactor are managed (this is discussed in greater detail in the section below on plutonium). Normally the term "plutonium breeder" in reserved for reactors specially designed and operated to produce high quality bomb-grade Pu.

Isotope production reactors use more highly enriched fuel, and heavy water as the moderator to minimize neutron losses. These reactors produce substantial excesses of neutrons that are absorbed by breeding materials inserted into the reactor. These reactors generally produce no power, but they are flexible for producing different isotopes. Lithium, depleted or natural uranium, or thorium can all be used. Since producing one atom of any isotope requires one neutron, the production rate for plutonium or U-233 in an isotope production reactor is 80 times higher than that of tritium.

The production in a breeder reactor is determined by the reactor operating power. Typically these reactors produce somewhat less than one atom of product for every atom consumed by fission. The breeding ratio (number of product atoms/number of fuel atoms) is usually 0.8-0.9. U.S. isotope reactors at Savannah, GA have a ratio of 0.86. A reactor consumes about one gram of fuel for every megawatt-day of operation,. A 100 megawatt reactor can thus produce about 85 g of Pu a day, or 1 g of tritium. Natural uranium fueled breeder reactors may have a higher breeding ratio than this for plutonium, but they have limited capacity for producing other isotopes.

The rate of isotope production can also be thought of as the product of the neutron flux intensity, the concentration of raw material in the reactor core, and the volume of the core. This view point helps explain another important factor, the build up of contaminating isotopes, which is a problem with plutonium and U-233 production.

Taking Pu-239 as an example, we note that initially there is no Pu-239 present in a freshly loaded reactor. Its concentration builds up at a rate proportional to: (neutron flux)*(U-238 concentration). This Pu-239 is essentially pure. The plutonium is exposed to the same flux as the U-238, so it also participates in neutron reactions. The rate of these reactions is proportional to the Pu concentration, so the rates are initially very small. Mostly these reactions are simply fission, partially offsetting the Pu production but also releasing more neutrons for breeding.

Pu-239 has a larger fission cross section than U-235, but it also has a much larger neutron absorption (non-fission capture) cross section. Absorption results in the following reaction:

Pu-239 + n -> Pu-240

As the concentration of Pu-239 grows, so too does the rate of Pu-240 production. Since the ratio of these two isotopes in the plutonium present in the reactor ultimately depends on the ratio of their respective production rates, the percentage of the plutonium consisting of Pu-240 grows steadily as more Pu accumulates.

Pu-240 in turn breeds Pu-241, which leads to Pu-242. The longer the irradiation of the U-238 continues (either in fuel rods, or in a Pu production blanket), the lower the percentage of the most desirable isotope (Pu-239) in the plutonium. Other plutonium isotopes are also produced (see below under Plutonium). Similar processes lead to contamination of reactor-produced U-233 as well.

6.2 Fissionable Materials

There are three isotopes known which are practical for use as fission explosives. These are U-235, Pu-239, and U-233. Of these only U-235 occurs in nature. Pu-239 and U-233 must be produced by bombarding other isotopes with neutrons. A third element, thorium (Th-232), can only undergo fast fission, but can also be used for breeding U-233. There are other elements that are also fissile but they have no practical significance for a variety of reasons. These elements are summarized in subsection 6.2.4.

6.2.1 Uranium (U)

Uranium, element 92, named after the planet Uranus, is usually given credit for being the last element on the periodic table to have survived from the time that the Earth formed (but see section 6.2.2 Plutonium below).

Its properties have been exploited at least since AD 79, since it has been found in yellow ceramic glazes (containing more than 1% uranium oxide) recovered from the ruins of Pompeii and Herculaneum, near Naples, Italy.

Uranium was discovered in 1789 in pitchblende by the German chemist Martin Heinrich Klaproth, who named it after the planet Uranus, which had been discovered in 1781. It was was first isolated as a metal in 1841 by French chemist Eugene-Melchior Peligot, who reduced the anhydrous uranium tetrachloride with potassium metal. A key milestone in the history of science was reached in 1896 when the French physicist Antoine Henri Becquerel discovered the existence of radioactivity by accidentally exposing a photographic plate by the ionizing radiations of potassium uranyl sulfate.

Uranium is a very heavy, silver-white lustrous metal. It is the eighth densest element. When pure it is a little softer than steel; is malleable, ductile, and slightly paramagnetic. Uranium has three allotropic forms, the alpha (orthorhombic, stable up to 667.7 C), the beta (tetragonal, stable from 667.7 C to 774.8 C), and the gamma (body-centered cubic, from 774.8 C to melting point) which is the most malleable and ductile. The alpha phase is very unusual type of

orthorhombic structure, consisting of corrugated sheets of atoms in a highly asymmetrical orthorhombic structure. This anisotropic, or distorted, structure makes it difficult for the atoms of alloying metals to substitute for uranium atoms or to occupy spaces between uranium atoms in the crystal lattice. Only molybdenum and niobium have been observed to form solid-solution alloys with uranium. Uranium metal does however react with a wide variety of alloying metals to form intermetallic compounds.

The basic properties of pure uranium are:

melting point 1132.2 degrees C (+/- 0.8); similar to that of copper

boiling point 3818 degrees C

density 18.95 (alpha form, stable to 667.8 C)

specific heat 6.65 cal/mole/C (25 C)

tensile strength 450 MPa

It is chemically very reactive. It oxidizes readily, a piece of metal progressively turns different colors as a thin oxide layer builds up (the same optical phenomenon that makes oil films colorful). Powdered uranium is pyrophoric, spontaneously bursting into flame in air. It ignites in air at 150-175 C forming U3O8. At 1000° C (1832° F), uranium combines with nitrogen to form a yellow nitride. Water attacks the metal, slowly when cold and faster when boiling. Uranium is soluble in hydrochloric, nitric and other acids forming tetravalent salts. It is not attacked by alkalies. Uranium displaces hydrogen from mineral acids and from the salt solutions of such metals as mercury, silver, copper, tin, platinum, and gold. On vigorous shaking metallic particles exhibit luminescence.

Uranium has four oxidation states, III-VI. The hexavalent compounds include uranyl trioxide, UO3, and uranyl chloride, UO2Cl2. Uranium tetrachloride, UCl4, and uranium dioxide, UO2, are examples of the tetravalent, or uranous, compounds. Uranous compounds are usually unstable; they revert to the hexavalent form when excessively exposed to air. Uranyl salts, such as uranyl chloride, may decompose in the presence of strong light and organic matter.

Uranium has fourteen isotopes, all radioactive. In nature only three are found. The approximate isotopic composition of natural uranium, with half-lifes is:

	Mass %	Atomic %	Half Life
U-234	0.0054%	0.0055%	247 thousand years (alpha emitter)
U-235	0.7110%	0.7202%	710 million years (alpha emitter)
U-238	99.2836%	99.2742%	4.51 billion years (alpha emitter)

Although the U-235 content of uranium is in general constant, occasional variations do occur on the order of 0.1% of the U-235 content (that is in the range of 0.7195-0.7209% as atomic %) due to depletion from ancient fission reactions that took place when U-235 existed in much higher concentrations than today. The 1.9 billion year old reactor centers discovered in 1972 at the Oklo mine in Gabon are the most famous. When the Oklo reactors were active natural uranium was about 3% U-235, the same as commercial power reactor fuel today. In the most highly depleted core at Oklo the U-235 content was only 0.44%. The Oklo reactor cores, and a few others found at another Gabon site near Oklo, are the only intact ancient reactors discovered so far.

The U-234 content can also vary significantly. Unlike U-235 and U-238 the isotope U-234 is not a ancient survival from stellar nucleosynthesis, it is an indirect decay product of U-238 and is produced by the decay chain:

U-238 -> (4.51 billion years, alpha) -> Th-234

Th-234 -> (24.1 day, beta) -> Pa-234

Pa-234 -> (6.75 hr, beta) -> U-234

Normally the U-234 exists in equilibrium with U-238, with the U-234 decaying away as fast as it is formed. Since the decaying atoms exist briefly as thorium and protactinium isotopes prior to the formation of U-234, they can become separated from the uranium bearing ore through chemical or physical processes. Ground water can preferentially leach one or both isotopes from uranium bearing soil, giving rise to ground water uranium with an isotopic composition enriched as much as 13 fold in U-234. Due to the short half-life, essentially all of the U-234 present in uranium has been formed in the last few million years. In equilibrium it accounts for about half of natural uranium's radioactivity.

U-236, with a half-life of 23.9 million years, is not found in nature in significant quantities. It accumulates when uranium is exposed to neutrons in reactors and can thus be used as a tracer for identifying uranium containing reprocessed fuel.

During the Manhattan Project natural uranium was code named "tuballoy" (abbreviated "Tu"), after the code named Tube Alloy Division of the project. This name is still encountered occasionally to refer to natural or depleted uranium. The code name for highly enriched uranium (specifically weapon-grade uranium) was "oralloy" (abbreviated "Oy"), which is also still encountered. The names "Q-metal", "depletalloy", and "D-38", once applied to depleted uranium, have fallen into disuse.

The specific activity of natural uranium is 0.67 microCi/g (split almost 50/50 between the rare U-234 isotope, and U-238; U-235 is a minor contributor). Natural uranium is sufficently radioative to expose a photographic plate in an hour or so.

A particularly important uranium compound is uranium hexafluoride (UF6). This is the only uranium compound that is stable and volatile near room temperature, and is used for gas phase uranium enrichment processes (diffusion and centrifugation). In this application it has the additional advantages that fluorine has only a single isotope (thus adding no complicating mass differences), and that UF6 is a stoichiometric compound (consisting of exactly 6 atoms of fluorine to 1 of uranium). It forms colorless crystals at room temperature, but sublimes (evaporates without melting) at 56 C. Its melting point is 64 C. It has a density of 4.87 (solid) and 3.86 (liquid). It is a powerful fluorinating agent, attacking most metals and oxides. It does not attack aluminum oxide (alumina). Aluminum metal resists attack (due to the thin alumina coating the metal possesses) as does nickel (due to the formation of a nickel fluoride layer). Most equipment for handling UF6 are made of aluminum or nickel alloys, or are nickel plated. Fluorinated hydrocarbons (e.g. Teflon) are excellent. It reacts instantly with water. Reactions with trace contaminants produce non-volatile lower fluorides that clog equipment.

As a footnote to uranium's role in weapons research, the compound uranium hydride, UH3, deserves mention. It was extensively studied at Los Alamos during the Manhattan Project as a possible bomb material. The theory was that the presence of hydrogen would moderate the neutrons, slowing the down to a region where U-235 has a much higher fission cross section. Although this would make the bomb less efficient, it was hoped that much smaller amounts of fissionable material might be needed. Further research showed that the unexpectedly low density (8), and the actual capture cross section of U-235 made this scheme unworkable. Implosion tests in the 1950s with (U-235)H3 cores verified this by producing negligible yields.

Before World War II uranium was thought to be a rare metal, it is now known to be more common than mercury or silver, and about as abundant as arsenic and molybdenum as a mineable ore. It has an average concentration in the crust of the earth of about 2 parts per 1 million, and, among the elements, ranks about 48th in natural abundance in crustal rocks. The discovery of large supplies has driven the price down to several dollars a pound. In the lithosphere (earth's crust) it is actually more common than such inexpensive elements as boron and zinc, existing in average quantities of 4 g/tonne. It is abundant enough in granitic rock to make radon gas (a decay product) a serious health hazard where outcrops are common. It is found in seawater at a concentration of 150 micrograms/tonne (cubic meter), the oceans are estimated to contain 4.5×10^9 tons of the element.

Uranium occurs as a significant constituent in more than 150 different minerals and as a minor component of another 50 minerals. Primary uranium minerals, found in magmatic hydrothermal veins and in pegmatites, include uraninite and pitchblende (the latter a variety of uraninite). The uranium in these two ores occurs in the form of uranium dioxide, which--owing to oxidation--can vary in exact chemical composition from UO_2 to $UO_{2.67}$. Other uranium ores of economic importance are autunite, a hydrated calcium uranyl phosphate; tobernite, a hydrated copper uranyl phosphate; coffinite, a black hydrated uranium silicate; and carnotite, a yellow hydrated potassium uranyl vanadate.

Uranium ores are widely distributed throughout the world. Reserves and commercial uranium transactions are expressed in U_3O_8 equivalent masses. The largest estimated reserves (May

1995, excluding Russia and Uzbekistan) are: Australia (27%), Kazakhstan (18%), Canada (14%), South Africa(9%), Namibia (9%), Brazil (8%), and the U.S. (4%). Russia and Uzbekistan also have large supplies, and commercial supplies have also been found in Gabon, Congo (or the Democratic Republic of the Congo, formerly Zaire), Kyrgyzstan, Czech Republic, Ukraine, and the UK (Cornwall). Worldwide reserves of inexpensive uranium ore (defined as uranium recoverable at a cost of U.S.$80/kg or less) is estimated at 3386 million kg U_3O_8 (again excluding Russia and Uzbekistan).

Deposits of pitchblende, the richest uranium ore, are found chiefly in Canada, Congo, and the United States. Most of the uranium mined in the United States is obtained from carnotite occurring in Colorado, Utah, New Mexico, Arizona, and Wyoming. A mineral called coffinite, discovered in 1955 in Colorado, is a high-grade ore containing nearly 61 percent uranium. Coffinite deposits were found subsequently in Wyoming and Arizona and in several foreign countries. In 1990, U.S. production of pure uranium concentrate was about 3417 metric tons. As of the end of 1994, known U.S. uranium reserves U_3O_8 were 133 million kg U_3O_8, and 432 million kg U_3O_8 at U.S.$110/kg. Approximately 73 percent of these reserves are located in deposits in New Mexico, Texas, and Wyoming. The average 1994 spot price for natural uranium in the United States in was $20.50/kg, a record low. Between 1985 and 1994 U.S. uranium prices averaged U.S.$69.51/kg U3O8 (U.S. suppliers) and U.S.$42.27/kg (foreign suppliers). World unrestricted prices were at U.S.$18.16 in October 1994, and are now (June 1996) at U.S.$38.90.

Uranium ores typically contain only a small amount of uranium-bearing minerals, so initial extraction and concentration steps are required. Physical separation steps (gravity, flotation, electrostatics, etc.) are not useful for uranium, so hydrometallurgical extraction (e.g. leaching) is the usual first step in ore processing.

In the classical acid leaching procedure for extracting uranium, the pitchblende ore is first broken up and roasted to dehydrate it, remove carbonaceous materials, oxidize sulfur compounds to sulfates, and oxidizes any other reductants that may interfere in subsequent leaching operations. It is then mixed with sulfuric and nitric acids. Uranium dissolves to form uranyl sulfate, UO_2SO_4; radium and other metals in the pitchblende ore are precipitated as sulfates. With the addition of sodium hydroxide, uranium is precipitated as sodium diuranate, Na2U2O7.6H2O, known as "yellow cake".

To obtain uranium from carnotite, the ore is finely ground and treated with a hot solution of caustic soda and potash to dissolve out uranium, radium, and vanadium. After the worthless sandy matrix is washed away, the solution is treated with sulfuric acid and barium chloride. A caustic alkali solution added to the remaining clear liquid precipitates the uranium and radium in concentrated form.

These classical methods of extracting uranium from its ores are now supplemented in current practice by such procedures as solvent extraction, ion exchange, and volatility methods.

In solvent extraction processes the uranium ore is removed from acidic ore leach-liquors by extracting with a solvent mixture such as tributylphosphate in kerosene. In modern industrial ore extraction methods, alkyl phosphoric acids--e.g., di(2-ethylhexyl) phosphoric acid--and secondary and tertiary alkyl amines are the usual solvents.

As a general rule, solvent extraction is preferred over ion-exchange methods for acidic leachates containing more than one gram of uranium per litre. Solvent extraction is not useful for recovery of uranium from carbonate leach liquors, however.

Uranium meeting nuclear-grade specifications is usually obtained from yellow cake through an additional solvent extraction step using the tributyl phosphate solvent-extraction process. First, the yellow cake is dissolved in nitric acid to prepare a feed solution. Uranium is then selectively extracted from this acid feed by tributyl phosphate diluted with kerosene or some other suitable hydrocarbon mixture. Finally, uranium is stripped from the tributyl phosphate extract into acidified water to yield a highly purified uranyl nitrate.

The nitrate is calcined to form UO_3, which is reduced in furnaces under a hydrogen atmosphere to UO_2. The UO_2 is converted to UF_4 with anhydrous HF.

The metal is usually prepared by reducing uranium halides (commonly the tetrafluoride) with magnesium by an exothermic reaction in a "bomb" (a refractory-lined, sealed container, usually steel), a general technique known as the Thermite process. The thermite reaction proceeds at temperatures exceeding 1,300 C. The production of uranium metal by magnesium reduction of the tetrafluoride specifically is also called the Ames process, after Iowa State University, Ames, where the American chemist F.H. Spedding and his colleagues developed it to industrial status in 1942. The bomb is necessary because the vapor pressure of magnesium metal is very high at 1300 C. The bombs are charged with granular UF_4 and finely divided Mg (the latter in excess) and are heated to 500 to 700 C, at which point an exothermic thermite reaction begins. The heat of reaction is sufficient to liquefy the conversion contents of the bomb, which are essentially metallic uranium and a slag of magnesium fluoride (MgF2). The slag separates and floats to the top. When the bomb is cooled to ambient temperature, a uranium metal ingot or "button" is obtained which, despite its hydrogen content, is the best-quality uranium metal available commercially and is well suited for rolling into fuel shapes for nuclear reactors.

The metal can also be prepared by reducing uranium oxides with calcium, aluminum, or carbon at high temperatures; or by electrolysis of KUF5 or UF_4, dissolved in a molten mixture of CaCl2 and NaCl. High-purity uranium can be prepared by the thermal decomposition of uranium halides on a hot filament.

It has been observed that the energy contained in coal in the form of trace uranium exceeds the energy content derived by combustion (assuming the U-238 is fissioned as well).

6.2.1.1 U-235

This relatively scarce isotope is the only naturally occurring material suitable for the production of energy through fission. For use in a fission weapon, or for convenient power production, it

needs to be concentrated to levels higher than that found in nature. Although civilian power plants require uranium with a concentration of 2-4.5% U-235, weapons require a minimum of 80% U-235, and preferably more than 90%. U.S. weapon grade uranium is about 93.5% U-235, U.S. enrichment plants are capable of producing a 97.65% "top product" (this is used in naval reactors). Uranium enriched to 80% or more is known generically as HEU, "highly enriched uranium" (a cut-off of 20% is also used). U.S. weapons grade uranium is called "oralloy", a wartime code-name derived from "Oak Ridge Alloy" that has remained in use. In 1998 ORNL Isotopes Division was offering weapon grade (93% U-235) for sale at $53/gram. Uranium with enrichments ranging from 40% to 80% U-235 has been used in large amounts in U.S. thermonuclear weapons as a yield-boosting jacketing material for the secondary fusion stage.

Since it is lighter still, U-234 is proportionately enriched even more than U-235 by mass-based enrichment processes. Highly enriched U-235 typically contains 1.5-2.0% U-234.

U-235 has a spontaneous fission rate of 0.16 fissions/sec-kg. A pure mass of U-235 weighing 60 kg would thus emit only 9.6 fissions/sec, making gun assembly quite easy. U-238 produces 35 times as many neutrons per kg, so even a small percentage percent of U-238 contaminant multiples this rate several fold. The U-234 contaminant with a fission rate 22 times higher has a similar effect. The specific activity for U-235 is 2.1 microCi/g (for the pure isotope); 0.8% of U-234 would increase this to 51 microCi/g.

6.2.1.1.1 U-235 Isotope Enrichment

Five technologies have been used for the large scale concentration of U-235:

- Electromagnetic separation
- Gaseous diffusion
- Thermal liquid diffusion
- Gas Centrifuge Separation
- Aerodynamic Separation

In addition serious attention continues to be paid to the following techniques:

- AVLIS
- Chemical separation

The first three technologies listed (electromagnetic, gaseous diffusion, and thermal liquid diffusion) were used during World War II by the U.S. to separate significant amounts of U-235.

The separative capacity of a plant (or a single stage) is measured in mass separative work units (SWU) per unit time, for example kg-SWU/yr of ton-SWU/yr. The output of enriched product from a plant of a given capacity is determined also by the concentration of isotope in the product, feed, and waste. The concentration in the feed is usually fixed by nature, but the feed and waste concentrations can be varied. Allowing a higher concentration in the waste (or "tails") boosts product output, but at the expense of more feed per unit of product. The relationship is:

$$P = U/[V(N_p) + [(N_p-N_f)/(N_f-N_w)]V(N_w) - (1+[(N_p-N_f)/(N_f-N_w)])V(N_f)]$$

where P is the product output rate, U is the separative capacity; N_p, N_f, N_w are the molar concentrations of isotope in the product, feed, and waste (measured in fractions, not percent); and $V(N_p)$, $V(N_w)$, $V(N_f)$ are the separative potential functions for each concentration. The separative potential function is defined as:

$V(N) = (2N-1)Ln(N/1-N)$ where Ln is the natural logarithm

Assuming tails of 0.25%, a plant with a capacity of 3100 kg-SWU/yr will produce 15 kg of 90% U-235 annually (enough for one implosion-type bomb) from natural uranium. It should be observed that if 3% U-235 is used as feed (commercial pressurized water reactor fuel), and natural concentration uranium is the waste, then only 886 kg-SWU/yr are required for this same output rate. Even low-enriched uranium could greatly assist a weapon production program if it was diverted.

Worldwide enrichment capacity is currently at 49 million separative work units (SWU). An annual demand of 26 to 38 million SWU is projected over the next 20 years.

Electromagnetic Separation

This was the first technique capable of producing bomb grade uranium to be perfected, and is probably the easiest for a low technology country to acquire. It was used in the Y-12 plant built at Oak Ridge by the U.S. during W.W.II. Y-12 used two separation stages to produce material of 80-90% purity. Low enrichment uranium tetrachloride feed produced by the other two processes was used to boost output compared to using natural uranium. All of the U-235 in the Hiroshima bomb was purified by this technique. This method was abandoned in 1946 due to its cost. Apparently the only country to pursue this method seriously since has been Iraq. Gaseous Diffusion

This was the first practical industrial scale uranium isotope separation technique. This method uses the different molecular speeds of molecules in uranium hexafluoride vapor to achieve separation. Due to the fact that thousands of stages are needed for high enrichment, a diffusion plant producing weapon-grade U-235 must be enormous and to be practical have a large production capacity as well. Developing suitable diffusion barriers that can resist UF6 corrosion is difficult. Barriers are of two general types, thin film tubes and aggregated particle tubes. Thin films use a thin barrier with diffusion pores formed by etching. Nitric acid leaching of 40/60 Au/Ag alloy or Ag/Zn alloy have been studied. Electrolytic etching of aluminum foil forms a brittle porous alumina barrier. Teflon films have also been studied. Aggregate barriers are made of small particles that are sintered or compacted to form relatively thick porous walls. The barrier material suitable for large scale production appears to have been sintered nickel powder. Diffusion barrier technology has been kept classified by all countries that have developed it.

A 3024 stage plant called K-25 was built at Oak Ridge by the U.S. during W.W.II, and continued operating until the 1970s. Developing suitable barrier materials was difficult and held up completion of the plant until after the end of the war, although the partially completed plant

contributed to the production of Little Boy. The barriers in use at this time were tubes made from sintered nickel powder, efforts to get a promising film barrier (probably electrolytically etched aluminum) into production in time failed. K-25 originally had 162,000 m^2 of barrier surface. This plant, with expansions, produced the majority of the fissionable material for U.S. weapons into the sixties. Improvements in barrier technology allowed performance enhancement by a factor of 23 by 1974.

Although basic plant construction is within the capabilities of standard chemical engineering technology today, developing barriers is a problem that must be overcome. Diffusion plants consume much less energy than electromagnetic plants, but the consumption is still quite large. Capital costs still dominate the cost of operation however. After an efficiency upgrade program, U.S. plants in 1981 had a specific power consumption of 2370 KWhr/kgSWU. Although low enriched uranium is valuable as feed to a plant intended to produce highly enriched uranium (HEU), a gaseous diffusion plant built specifically for low enriched uranium production cannot be easily diverted to produce HEU. High enrichment calls for diffusion stages much smaller than the large units used for low enriched material, both because the feed rate is dramatically lower, and because of criticality problems if a stage hold too much material. A commercial gaseous diffusion plant typically takes 1-2 months to reach equilibrium after start-up (an HEU plant would take longer, due to the greater number of successive stages).

A number of countries have used this technology, including Argentina which successfully developed gaseous diffusion technology and operated a plant in a secret weapon program (now terminated). In 1979 more than 98% of all uranium separation capacity in the world used this process. By the mid 80s, this had dropped to 95% as new gas centrifuge plants came on line.

Thermal Liquid Diffusion

The first technology to produce substantial amounts of (slightly) enriched uranium. This technology was used by the U.S. during W.W.II to boost the output of the Y-12 plant by providing slightly enriched feedstock. The thermal diffusion technology is the simplest of all these techniques, but is capable of enriching U-235 only to around 1% (the S-50 plant at Oak Ridge produced 0.85-0.89% U-235). Immense amounts of heat are required.

Gas Centrifuge Separation

Although Germany developed this technology during W.W.II, it did not see application until the sixties. It is the dominant U-235 separation technique for new plant construction today, although existing capacity is still mostly gaseous diffusion. By passing uranium hexafluoride vapor through high speed centrifuges, the isotopes concentrate in lighter and heavier gas layers which are bled off. Each centrifuge produces a much larger separation factor than a gaseous diffusion barrier. Fewer separation stages are needed, on the order of a thousand, but the cost of each stage is much higher than for a gaseous diffusion plant. Gas centrifuge separation requires less than 10% of the power required by gaseous diffusion, and it is easier to adjust the scale of production plants. A modern gas centrifuge provides at least 6 kg-SWU/yr of separative capacity. The specific power consumption is 100-250 kWhr/kgSWU. A high level of technical

proficiency is required to build efficient centrifuges. This technology was successfully acquired by Pakistan and India, and was avidly pursued by Iraq and Brazil.

Aerodynamic Separation

This technology was developed by South Africa (UCOR process, using vortex tubes operating at a pressure of 6 bar) and Germany (using curved nozzles operating below atmospheric pressure, 0.25-0.5 bar). South Africa has been the only country to use it on a large scale. It produced about 400 kg of bomb-grade uranium at a plant at Valindaba, which was shut down in the late eighties. Separation factors of at least 1.015 have been achieved. It requires as much or more energy as gaseous diffusion, a typical figure is 3300 kWhr/kgSWU. AVLIS (atomic vapor laser isotope separation)

Never used in practice, this technology was being developed in the U.S. during the 70s and 80s until it was killed due to a world-wide glut of isotope separation capacity and reductions in the strategic arsenal. A billion dollar AVLIS plant was planned for Idaho Falls, Idaho during the eighties but was never constructed. Chemical Separation

Chemical exchange techniques for uranium have also been developed by Japan and France, but have never been used for a production plant. The French Chemex method involves a countercurrent of two immiscible liquids, each containing dissolved uranium, in a tall column. The Japanese Asahi process uses an exchange reaction between an aqueous solution and a finely powdered resin through which the solution percolates slowly. Both require special catalysts to accelerate the concentration process. The Chemex process has demonstrated a power consumption of 600 kWhr/kgSWU. Iraq pursued this technology (in the form of a combined Chemex/Asahi plant) as a means to partially enrich U-235 to 6-8% for use as a calutron feedstock. By cutting the amount of feed to be processed, output would have been boosted 8-11 times.

Approximate energy efficiency ranking of uranium enrichment techniques is given below, with typical relative power consumptions (gaseous diffusion = 1):

less than 0.01? AVLIS (if developed to production status)

0.10-0.04	Gas Centrifuge
0.30	Chemical Separation
1.00	Gaseous Diffusion
1.50	Aerodynamic Enrichment
high	Electromagnetic Separation
high	Thermal Liquid Diffusion

6.2.1.2 U-238

Although it can't be used as a primary fissile material due to the high threshold energy required for neutrons to cause fission, U-238 is still a very important nuclear material.

Due to its high density and atomic weight it is useful in bomb construction as a tamper/reflector for both fission and fusion assemblies. The fact that it *is* fissioned by fast neutrons means that it can boost weapon yield indirectly by multiplying reflected neutrons, or directly by the fast fission of the tamper. About 40% of fission neutrons, and all fusion neutrons, are energetic enough to fission U-238.

U-238 has a spontaneous fission rate some 35 times that of U-235, 5.51 fissions/sec-kg. This prevents its use as a tamper/reflector in bombs using gun assembly since a reasonable amount (200-300 kg)would emit too many neutrons. Pure U-238 has a specific activity of 0.333 microCi/g.

It is also the raw material for breeding plutonium-239, a very valuable primary fissile isotope. Any reactor that uses natural or partially enriched uranium (i.e. nearly every reactor in the world) contains U-238 and thus breeds plutonium.

6.2.1.3 U-233

This fissile uranium isotope (half-life 162,000 years) is not found in nature. It is instead bred from thorium-232 in a manner similar to the production of Pu-239:

Th-232 + n -> Th-233

Th-233 -> (22.2 min, beta) -> Pa-233

Pa-233 -> (27.0 day, beta) -> U-233

A two-step side reaction chain also occurs during breeding leading to the production of U-232:

Th-232 + n -> Th-231 + 2n

Th-231 -> (25.5 hr, beta) -> Pa-231

Pa-231 + n -> Pa-232

Pa-232 -> (1.31 day, beta) -> U-232

The production of U-232 through this process depends on the presence of significant amounts of un-thermalized neutrons since the cross section of the initial n,2n reaction is small at thermal energies.

If significant amounts of the isotope Th-230 are present then U-232 production is augmented by the reaction:

Th-230 + n -> Th-231

which continues as before.

The presence of U-232 is important because of its decay chain:

U-232 -> (76 yr, alpha) -> Th-228

Th-228 -> (1.913 yr, alpha) -> Ra-224

Ra-224 -> (3.64 day, alpha & gamma) -> Rn-220

Rn-220 -> (55.6 sec, alpha) -> Po-216

Po-216 -> (0.155 sec, alpha) -> Pb-212

Pb-212 -> (10.64 hr, beta & gamma) -> Bi-212

Bi-212 -> (60.6 min, beta & gamma) -> Po-212

alpha & gamma) -> Tl-208

Po-212 -> (3x10^-7 sec, alpha) -> Pb-208 (stable)

Tl-208 -> (3.06 min, beta & gamma) -> Pb-208

The rapid decay sequence beginning with Ra-224 produces a large amount of energetic gamma rays. About 85% of this total gamma energy output is due to the last isotope in the sequence, thallium-208 which produces the most energetic gamma rays (up to 2.6 MeV). The amount of gamma radiation emitted is proportional to the amount of Th-228 present.

The buildup of U-232 as a contaminant is unavoidable during the production of U-233. This is similar to the plutonium isotope contamination problem discussed below in plutonium

production, but occurs to a much smaller extent rate. The first (n,2n) reaction only occurs when neutrons with energies in excess of 6 MeV are encountered. Only a small percentage of fission neutrons are this energetic, and if the thorium breeding blanket is kept in a reactor region where it is only exposed to a well moderated neutron flux (i.e essentially no neutrons above the Th-232 fission threshold of 500 KeV) this reaction can be nearly eliminated. The second reaction proceeds very efficiently with thermalized neutrons however, and minimizing U-232 from this source requires choosing thorium that naturally has a low Th-230 concentration.

If the above precautions are followed weapons-grade U-233 can be produced with U-232 levels of around 5 parts per million (0.0005%). Above 50 ppm (0.005%) of U-232 is considered low grade.

In a commercial fuel cycle the build-up of U-232 is not really a disadvantage, and may even be desirable since it reduces the proliferation potential of the uranium. In a fuel economy where the fuel is reprocessed and recycled the U-232 level could build up to 1000 - 2000 ppm (0.1 - 0.2%). In a system that is specifically engineered to accumulate U-232 levels of 0.5-1.0% can be reached.

Over the first couple years after U-233 containing U-232 is processed, Th-228 builds up to a nearly constant level, balanced by its own decay. During this time the gamma emissions build up and then stabilize. Thus over a few years a fabricated mass of U-233 can build up significant gamma emissions. A 10 kg sphere of weapons grade U-233 (5 ppm U-232) could be expected to reach 11 millirem/hr at 1 meter after 1 month, 0.11 rem/hr after 1 year, and 0.20 rem/hr after 2 years. Glove-box handling of such components, as is typical of weapons assembly and disassembly work, would quickly create worker safety problems. An annual 5 rem exposure limit would be exceeded with less than 25 hours of assembly work if 2-year old U-233 were used. Even 1 month old material would require limiting assembly duties to less than 10 hours per week.

In a fully assembled weapon exposures would be reduced by absorption by the tamper, case, and other materials. In a modern light weight design this absorption would be unlikely to achieve more than a factor of 10 attenuation, making exposure to weapons assembled two years previously an occupational safety problem. The beryllium reflectors used in light weight weapons would also add to the background neutron level due to the Be-9 + gamma -> Be-8 + neutron reaction. The U-232 gammas also provide a distinctive signature that can be used to detect and track the weapons from a distance. The heavy tampers used in less sophisticated weapon designs can provide much high levels of attenuation - a factor of 100 or even 1000.

With deliberately denatured grades of U-233 produced by a thorium fuel cycle (0.5 - 1.0% U-232), very high gamma exposures would result. A 10 kg sphere of this material could be expected to reach 11 rem/hr at 1 meter after 1 month, 110 rem/hr after 1 year, and 200 rem/hr after 2 years. Handling and fabrication of such material would have to done remotely (this also true of fuel element fabrication) In an assembled weapon, even if a factor of 1000 attenuation is assumed, close contact of no more than 25 hours/year with such a weapon would be possible

and remain within safety standards. This makes the diversion of such material for weapons use extremely undesirable.

The short half-life of U-232 also gives it very high alpha activity. Denatured U-233 containing 1% U-232 content has three times the alpha activity of weapon-grade plutonium, and a correspondingly higher radiotoxicity. This high alpha activity also gives rise to an even more serious neutron emission problem than the gamma/beryllium reaction mentioned above. Alpha particles interact with light element contaminants in the fissile material to produce neutrons. This process is a much less prolific generator of neutrons in uranium metal than the spontaneous fission of the Pu-240 contaminant in plutonium though.

To minimize this problem the presence of light elements (especially, beryllium, boron, fluorine, and lithium) must be kept low. This is not really a problem for U-233 used in implosion systems since the neutron background problem is smaller than that of plutonium. For gun-type bombs the required purity level for these elements is on the order of 1 part per million. Although achieving such purity is not a trivial task, it is certainly achievable with standard chemical purification techniques. The ability of the semiconductor industry to prepare silicon in bulk with a purity of better than one part per billion raises the possibility of virtually eliminating neutron emissions by sufficient purification.

U-233 has a spontaneous fission rate of 0.47 fissions/sec-kg. U-232 has a spontaneous fission rate of 720 fissions-sec/kg.

Despite the gamma and neutron emission drawbacks, U-233 is otherwise an excellent primary fissile material. It has a much smaller critical mass than U-235, and its nuclear characteristics are similar to plutonium. The U.S. conducted its first test of a U-233 bomb core in Teapot MET in 1957 and has conducted quite a number of bomb tests using this isotope, although the purpose of these tests is not clear. India is believed to have produced U-233 as part of its weapons research and development, and officially includes U-233 breeding as part of its nuclear power program.

Its specific activity (not counting U-232 contamination) is 9.636 milliCi/g, giving it an alpha activity (and radiotoxicity) about 15% of plutonium. A 1% U-232 content would raise this to 212 milliCi/g.

6.2.1.4 Depleted Uranium

When U-235 has been extracted from natural uranium, the remaining uranium is called "depleted uranium" because it has been depleted of U-235. Typically it still contains 0.25-0.4% U-235 since it is not economical to strip U-235 in concentrations this low (it is cheaper just to buy more natural uranium with higher concentrations). The U.S. tails concentration was 0.2531% in 1963, rising to 0.30% in the 70s. Due to the recent low cost of uranium, the trend has been toward higher tails concentrations in recent years.

The U.S. has some 560,000 tonnes of depleted uranium hexafluoride (UF6) currently stored in cylinders at DOE's three gaseous diffusion plant sites: Paducah, Kentucky; Portsmouth, Ohio; and Oak Ridge, Tennessee. (This is probably not the complete U.S. DU inventory.)

Depleted uranium is about half as radioactive as natural uranium, almost entirely due to the removal of U-234, but its bulk properties are otherwise identical. Because the major use of uranium is for power production, depleted uranium is a waste product with little economic value. Finding uses for depleted uranium stockpiles has been a problem for enrichment plants.

The uses it has found have generally been related to its great density and comparatively low cost. It two most important uses are as a radiation shielding material, and as balance weights in aerospace applications such as aircraft control surfaces, each Boeing 747 contains 1500 kg of DU for this purpose. It is used extensively in oil-well drilling in the form of sinker bars, weights used to lower instruments into mud-filled well holes. It has also been used in high-performance gyroscope rotors, large inertial energy-storage flywheels, and as ballast in ballastic missile reentry vehicles and racing yachts.

Most famously, it is used in armor-piercing ammunition by the U.S.. When suitably alloyed and heat treated (alloyed with 2% Mo or 0.75% Ti; rapid quenching for 850 C in oil or water, followed by aging at 450 deg. for 5 hrs) uranium is as hard and strong as hardened tool steel (tensile strength >1600 MPa). Combined with its great density this makes it very effective at punching through armor, consistently superior to the much more expensive monocrystalline tungsten, it nearest competitor. The penetration process pulverizes much of the uranium, which explodes into burning fragments when it hits the air on the other side, adding to the destruction. Some 300 tonnes of depleted uranium were fired on the battlefield during Desert Storm (mostly from the GAU-8 30-mm guns of A-10 attack planes, each round contains 272 grams of DU alloy). It is also used in advanced armor, like that used on the M-1 Abrams tank.

6.2.2 Plutonium (Pu) Plutonium, element 94, named after the planet Pluto, was discovered by Glenn Seaborg, Edwin McMillan, Kennedy, and Arthur Wahl in 1940 at Berkeley by bombarding uranium with deuterons from the 60 inch cyclotron.The fissile properties of Pu-239 had been predicted in advance of the element's discovery by Louis Turner in May 1940. The isotope Pu-239 exists naturally in trace amounts in uranium ore (several parts per quadrillion) produced by the capture of spontaneous fission neutrons by U-238. Extremely small amounts of Pu-244 (the longest lived plutonium isotope, with a half-life of 80 million years) have been reportedly detected in cerium ore, apparently surviving remnants of plutonium present at the formation of the earth.

Plutonium is produced by bombarding U-238 with slow neutrons, in a nuclear reactor for large amounts (a process called breeding). If a slow neutron is captured, then U-239 is produced which quickly decays into neptunium-239 and then plutonium:

U-238 + n -> U-239

U-239 -> (23.5 min., beta) -> Np-239

Np-239 -> (2.35 days, beta) -> Pu-239

Fifteen isotopes of plutonium are known, all are radioactive. The principle ones of interest to the design of nuclear weapons, and their half-lifes, are:

Pu-238 -> (86 years, alpha) -> U-234

Pu-239 -> (24360 years, alpha) -> U-235

Pu-240 -> (6580 years, alpha) -> U-236

Pu-241 -> (14.0 years, beta) -> Am-241

Pu-242 -> (370000 years, alpha) -> U-238

Plutonium is a very heavy silvery metal, bright like nickel when freshly cleaned. It is a highly electronegative, chemically reactive element, considerably more so than uranium. It quickly tarnishes, forming a succession of interference colors (like an oil film), initially a pale yellow, and eventually becoming dark purple-black. If corrosion is sufficiently severe, a loose olive-green powder (PuO_2) forms on the surface.

Plutonium oxidizes readily, and corrodes very quickly in the presence of even traces of moisture. Oddly, it corrodes much faster in an inert atmosphere if moisture is present, than it does in ordianry air or pure oxygen. The reason is that direct attack by oxygen forms an adherent plutonium oxide layer that retards further corrosion, attack by moisture produces a loose oxide and hydride mixture. efficient desiccator is necessary to preve and oxidizes and corrodes readily, and is attacked by moisture.

It has four valences, III-VI. It has significant solubility only in very acidic solutions, such as nitric and hydrochloric acid (as the nitrate and chloride), as well as the hydroiodic and perchloric acids. Plutonium salts hydrolyze readily in contact with neutral of basic solution, forming insoluble plutonium hydroxide. Concentrated plutonium solutions are unstable due to radiolytic decomposition, resulting in precipitation.

Due to radioactive heating, a substantial piece of plutonium is warm to the touch. A large piece of plutonium that is thermally insulated can exceed the boiling point of water.

The basic properties of plutonium are:

melting point 641 degrees C; similar to that of aluminum

boiling point 3232 degrees C

density 19.84 (alpha phase)

Plutonium has a large number of very peculiar properties. It has the lowest thermal conductivity of any metal, its electrical conductivity is lower than any metal but manganese (other reports claim flatly that it is the lowest). It is the most viscous liquid metal known. And it undergoes the most extreme and bizarre density changes with temperature of any element.

Plutonium has six distinct phases (crystal structures) in solid form, more than any other element (actually there is a seventh that exists under certain conditions). Some of the phase changes involve dramatic changes in volume. In two of these phases (delta and delta prime) plutonium has the unique property of shrinking with increasing temperature, in the others it has an extremely large thermal expansion coefficient. Plutonium also contracts upon melting, allowing unmelted plutonium to float. In its densest form, the alpha phase, plutonium is the sixth densest element (only osmium, iridium, platinum, rhenium, and neptunium are denser).

It is brittle in the alpha phase, which exists at room temperature in pure plutonium, but ductile alloys exist.

The densities and temperature ranges of these phases are:

Phase	Density	Stability Range (degrees C)
alpha	19.84 (20 C)	stable below 122
beta	17.8 (122 C)	122 - 206
gamma	17.2 (206 C)	224 - 300
delta/		
delta prime	15.9 (319 C)	319 - 476
epsilon	17.0 (476 C)	476 - 641 (melting point)
liquid	16.65 (641 C)	641 - to boiling point

By the end of 1995 about 1270 tonnes of plutonium had been produced world wide, 257 tonnes for weapon use, the rest as a by-product of commercial power production. Due to the rapid recent growth of the world wide nuclear power industry, reactor grade uranium is currently being generated by commerical operations at a rate of 75 tonnes/yr. About 210 tonnes of the commercial plutonium has been separated from nuclear fuel (and all of the military plutonium, of

course) so far. Reprocessing is producing about 18 tonnes a year right now (only the UK, France, Russia, Japan, and India engage in reprocessing).

Besides its use in nuclear weapons, plutonium has potential utility as a source of commercial power. Japan currently has the only major program for using plutonium for power. It appears that plutonium will not be economically competitive with uranium for decades to come for several reasons. The cost of reprocessing reactor fuel to extract plutonium is more expensive than the current commercial cost of low enriched uranium. Most current fuel fabrication plants are not equipped to handle the more hazardous plutonium oxide. The costs of safeguarding plutonium to prevent its theft or diversion for weapons use is quite substantial. Existing power reactors can only use fuel containing fairly small amounts of plutonium, which makes its use of little value, and the costs of designing and building new reactors for this purpose also adversely affects its cost. The current abundant supplies of uranium, excess enrichment capacity, and the large stockpiles of U.S. and Russian weapons grade uranium now being mixed down for commercial use guarantee rock bottom uranium prices for the next 20-30 years.

Plutonium has few other uses. The most widespread is its use as a smoke detector radioisotope in Europe (U.S. smoke detectors use americium due to its shorter half-life). Plutonium-beryllium alloy is used as a laboratory neutron source. The isotope Pu-238 is useful for radioisotope-powered thermoelectric generators (RTGs) to power long duration deep space probes due to its high heat output and long lifetime.

In Sept. 1998 the prices charged by the ORNL Isotopes Division for different isotopes of plutonium was $8.25/mg of Plutonium-238 (97% purity); $4.65/mg of Plutonium-239 (>99.99%); $5.45/mg of Plutonium-240 (>95%); $14.70/mg of Plutonium-241 (>93%); and $19.75/mg of Plutonium-242.

Although extremely scarce naturally, about 5000 kg of plutonium has been released into the atmosphere by nuclear weapons tests. The soil of the U.S. contains an average of about 2 millicuries (28 milligrams) per km^2 from fallout.

6.2.2.1 Plutonium Metallurgy

At room temperature pure plutonium exhibits a crystal structure that is termed the "alpha phase". In this form plutonium is at its maximum density, some 19.84 at 20 C. The atoms in the alpha phase are essentially covalently bonded (as opposed to metallically bonded), giving it physical properties more nearly like a mineral than a metal. It is hard, strong, brittle, and fractures in a highly directional manner. The alpha phase is unworkable using normal metal fabrication techniques.

In its lowest density phase (density of about 15.9), the delta phase, it is quite malleable. Plutonium is also quite malleable in the gamma phase.

Delta phase plutonium has normal metallic properties, including excellent ductility. The delta phase has strength and malleability similar to aluminum, which makes forming and machining

straightforward. Although the delta phase exhibits the anomalous property of shrinking when heated, this negative coefficient of expansion is not large.

The delta phase is only marginally stable. It tends to collapse into the dense alpha phase under very low pressures (as these things go), causing a 25% density increase. The delta phase does not exist in pure plutonium above a pressure of about 1 kilobar. For comparison, a density increase of 25% in uranium (or alpha phase plutonium) requires pressures of 450 kilobars. Above 30 kilobars only the alpha and beta phases of plutonium exist.

Plutonium can be stabilized in the delta phase at room temperature by alloying it with certain trivalent atoms such as gallium, aluminum, cerium, indium, scandium, and americium at concentrations of a few molar% (% of atoms that are the alloying agent).

Even when stabilized, the delta phase still collapses easily, a pressure of a few kilobars is sufficient to make it revert to the alpha phase. Of special interest is the fact that in gallium stabilized plutonium, the delta phase is actually metastable when the gallium content is below 4.0 molar%. This means that the pressure induced phase change to the alpha form is irreversible.

For use in weapons plutonium is normally stabilized in the delta phase by alloying it with 3-3.5 molar% gallium (0.9-1.0% by weight). This alloy is stable from at least -75 to 475 degrees C. The stabilization prevents low temperature phase changes from occurring after fabrication that would ruin the precisely made components of a weapon. It has an almost zero coefficient of expansion. It also makes casting easier since only the epsilon -> delta phase change occurs during cooling. Finally, the gallium reduces the susceptibility to corrosion of plutonium.

The 3% gallium alloy was used in the Gadget and Fat Man bombs. Aside from this alloying agent the plutonium was otherwise highly pure.

Alpha phase plutonium is known to have been used as weapon components however.

Aluminum is a good alloying agent, but early in the U.S. weapons program it was not used in weapons due to the existence of neutron producing alpha->n reactions (see below). Cerium is not used since (among other reasons) it does not confer corrosion resistance.

Plutonium pits are plated with metal (usually nickel) to protect them from corrosion, as well as reducing the radiological hazard. The hemispheres for Gadget were electroplated with nickel (also reported - apparently incorrectly - as being silver), a process that was not very successful and led to blistering. A process was developed for chemically plating the surface by exposing the plutonium pit to a nickel carbonyl atmosphere. The pits for the Fat Man bomb as well as the Operation Crossroads devices were plated with nickel, as was the pit for Joe 1/RDS-1. Evaporation coating with aluminum and electroplating with zinc do not work.

A potentially serious problem with using plutonium in weapons is the existence of a high spontaneous neutron emission rate. The presence of neutrons during the assembly of a supercritical mass leads to a premature nuclear reaction, an inefficient release of energy, even

a near total weapon failure in some cases. There are two sources for this neutron background. The most important is the presence of the isotope Pu-240, which spontaneously fissions fast enough to release some 10^6 n/sec/kg. This isotope inevitably forms during Pu-239 production. The second source is from the interaction of the strong alpha emissions with light element contaminants in the plutonium. Although this problem aroused great concern on the Manhattan Project when the use gun assembly was originally planned, the discovery of the Pu-240 problem rendered it moot. To minimize (but not eliminate) this problem the presence of light elements (especially beryllium, fluorine, and boron) must be kept to parts-per-million levels, a task of substantial difficulty. Aluminum also undergoes the alpha -> n reaction to some extent making it less desirable as an alloying agent in weapons, although with modern weapons grade plutonium this contribution to neutron emissions is insignificant.

Ultimately the satisfactory performance of gallium alloys, the established base of experience with them, and the relative unimportance of such issues as the cost of the alloying agent preclude the consideration of agents like aluminum.

The original techniques for preparing plutonium metal involved pyrochemical reduction of plutonium halides with alkali metals. Typically PuF4 was reduced with calcium and iodine, this was the standard method in the U.S. at least into the 1970s. Higher purity can be achieved by electrorefining of the pyrochemically produced metal (a step not necessary for weapons use). This makes use of an electrolytic cell at 700 C, with a sodium, potassium, and plutonium chloride electrolyte, and a tungsten or tantalum cathode, and produces 99.99% pure plutonium. More recent techniques are based on direct pyrochemical reduction and electrorefining of plutonium oxides. Among other advantages, these processes produce fewer waste products that must be disposed of.

Handling of molten plutonium, and plutonium casting, is performed today using equipment made of slightly oxidized tantalum. Casting molds can also be made of machined graphite, mild steel, or cast iron if they are lined with calcium fluoride or the oxides of zirconium or yttrium. It has also been discovered that pure plutonium can be successfully cast in chilled aluminum molds. The cooling is so fast that the intermediate phase transformations, which occur relatively slowly, are almost entirely bypassed.

6.2.2.2 Plutonium Toxicity

Although plutonium presumably exhibits chemical toxicity like other heavy metals, this effect is insignificant (in fact, unobservable) compared to its radiotoxicity. Plutonium's toxic properties are due to the fact that it is an active alpha emitter. Alpha particles are hazardous only if they are emitted inside the body (i.e. the plutonium has been ingested). Although plutonium emits gammas and neutrons that penetrate the body from outside, the emission rate is too small to be a significant hazard. Since the toxicity of plutonium is determined by its alpha activity, the isotopic composition is a significant influence. Weapon grade plutonium, free of americium, is assumed below (activity approx. 0.071 Ci/g).

Alpha particles affect only tissues that contain plutonium or are in direct contact with it. Two types of effects are significant: acute and chronic toxicity. If the exposure rate is high enough, tissues can suffer acute radiation poisioning with toxic effects appearing quickly. If the rate is low, then the cumulative carcinogenic effect occurs.

It is very poorly absorbed through the GI tract, even when ingested as a soluble salt since it tends to bind with the contents of the stomache and intestines. Given its tendency to precipitate from aqueous solution, and to form insoluble complexes with other materials, plutonium contamination of water tends to be a self-limiting phenomenon.

Swallowing 500 mg (7 curies) of plutonium as a finely divided or soluble material can cause death from the acute exposure of the GI tract in several days to a few weeks. Inhalation of 100 mg (1.4 Ci) of plutonium as particles of optimal size for lung retention can cause death from lung edema in 1 to 10 days. An inhaled dose of 20 mg (0.28 Ci) will cause death by fibrosis in about 1 month. In doses much below these values, the chronic carcinogenic effects become the important ones.

To exert chronic effects, the plutonium must be continuosly present in the body. Inhaled insoluble particles of the appropriate size range for lung retention (1-3 microns) will most likely be permanently deposited in the lungs (a high explosive detonation, like a non-nuclear weapons accident, can convert 20-50% of the plutonium present to this form). The most likely chemical form to which a person might be exposed is plutonium oxide. The oxide is used in reactor fuel, and metallic plutonium particles are rapidly oxidized. The oxide is nearly insoluble in water.

The lifetime risk of lung cancer from deposited plutonium particles for an adult is roughly proportional to the amount ingested. The ingestion of 1 microgram of Pu (0.07 microCi) represents a risk of 1% of developing cancer (the normal rate of cancer incidence is 20%). Thus ingesting 10 micrograms boosts the lifetime risk of cancer from 20% to 30%. Ingesting 100 micrograms (7.1 microCi) or more virtually guarantees eventual development of lung cancer (usually after several decades)although evidence of lung damage may surface within several months.

Plutonium normally exists in biological systems in the +4 oxidation state which chemically resembles Fe 3+. If it is absorbed into the circulatory system it thus a high probability of being concentrated in tissues that contain iron: bone marrow (which is highly sensitive to radiation), liver, and spleen. Plutonium has a biological half-life of 80-100 years when deposited in bone tissue, essentially permanent. Its biological half-life in the liver is 40 years. Chelating agents may help accelerate plutonium removal. If 1.4 micrograms (0.1 microCi) is deposited in an adult's bones, immune system impairment will result, and bone cancer is likely to develop within several years.

The International Commission on Radiological Protection (ICRP) specifies an Annual Limit on Intake (ALI) of 20 nanoCi/yr (280 nanograms). This translates into an air concentration of 7 picroCi/M^3 for occupational exposure. The maximum allowable body burden of Pu-239 (occupational exposure) is 40 nanoCi (0.56 micrograms) and 16 nanoCi for lung burden (0.23

micrograms). Due to the problems of contamination, skin contact with plutonium is strictly forbidden in U.S. laboratories.

6.2.2.3 Plutonium Production

Plutonium-239 is the preferred isotope for weapons use. As I have discussed above in the section on transmutation, it is produced in nuclear reactors where U-238 is exposed to a flux of slow neutrons. This occurs automatically in the vast majority of the world's reactors since they use low-enriched or natural uranium as fuel, which consists mostly of U-238. It can also be produced in special reactors that use highly enriched uranium or plutonium as fuel, but include a blanket of natural or depleted uranium for plutonium breeding.

Non-fission capture of neutrons by Pu-239 causes Pu-240 to form as irradiation proceeds. In turn, Pu-241 and Pu-242 also accumulate though in diminishing quantities.

Pu-239 + n -> Pu-240

Pu-240 + n -> Pu-241

Pu-241 + n -> Pu-242

A side reaction chain also produces Pu-238:

U-238 + n -> U-237 + 2n

U-237 -> (6.75 days, beta) -> Np-237

Np-237 + n -> Np-238

Np-238 -> (2.1 days, beta) -> Pu-238

The total irradiation that a fuel element or blanket element receives is measured in megawatt-days/tonne. When fuel elements are being discussed, this is referred to as the fuel's "burn-up". Better quality plutonium (for weapons) comes from elements (or "targets") with low MWD/tonne exposures which result in lower concentrations of the other undesired isotopes. The fuel elements in a modern enriched uranium light water reactor can reach 33000 MWD/tonne. Typical exposures in weapon breeder reactors are around 1000 MWD/tonne. Weapons grade plutonium in the Hanford graphite piles was produced with exposures of some 600 MWD/tonne, the Savannah River heavy water piles produced equivalent compositions with 1000 MWD/tonne

(presumably because part of the neutron flux is being absorbed by tritium production). During the Manhattan Project, the natural uranium fuel in the Hanford production reactors received only 100 MWD/tonne exposures, due to the urgency of wartime production, and produced super-grade weapons plutonium (0.9-1.0% Pu-240, and negligible amounts of other isotopes).

6.2.2.4 Pu-238

This isotope has a spontaneous fission rate, 1.1x10^6 fission/sec-kg (2.6 times that of Pu-240) and a very high heat output (567 W/kg!). Its very high alpha activity (283 times higher than Pu-239) makes it a much more serious source of neutron emission from the alpha -> n reaction. In high-burnup commerical reator fuels it makes up no more tha one or two percent of plutonium composition in extracted plutonium, but even so the neutron production and heating can make it very troublesome. It is used in radioisotope thermal generators (RTGs) which produce electricity for applications such as long duration space missions such as Cassini and deep sea intelligence gathering systems. For these purposes it is produced by bombarding pure neptunium-237 targets in breeding reactors. Its specific activity is 17.5 Ci/g.

6.2.2.5 Pu-239

Pu-239 is the only desired isotope for weapons use, other isotopes are important through their adverse effects. Pu-239 has higher fission and scattering cross sections than U-235, and a larger number of neutrons produced per fission, and consequently a smaller critical mass.

Pure Pu-239 has a moderate rate of neutron emission from spontaneous fission, about 10 fission/sec-kg (some 30 neutrons/sec-kg). Considering the small critical mass required, 6 kg or less, gun assembly could be used if pure Pu-239 were available (although, due to its high alpha activity, light element impurities would have to be kept to a few ppm to avoid alpha -> n reactions).

The relatively short half-life of Pu-239 (compared with U-235) means that a significant amount of energy is emitted through radioactive decay. In fact Pu-239 produces 1.92 watt/kg. This is higher than the average metabolic rate of an adult human by weight, and the output is concentrated in one-twentieth the volume. Pieces of Pu-239 are consequently quite warm. If a piece were thoroughly insulated, its temperature would rise from room temperature to the boiling point of water in less than two hours, and to the alpha-beta transition point soon after. This presents a problem in weapon design since elevated temperatures can be reached from self-heating, even if environmental heating is avoided. It is this concern about an alpha phase pit reaching the alpha->beta transition temperature that prevents this phase from being used in weapons. The specific activity of Pu-239 is 61.5 milliCi/g.

The plutonium used in the U.S. weapons program was produced at Hanford, Washington and Savannah River, Georgia. Soviet/Russian plutonium production is located at Kyshtym, near Chelyabinsk.

6.2.2.6 Pu-240

Pu-240 is the major contaminant of concern in plutonium intended for weapons use. The level of isotopic contamination by Pu-240 is important primarily for its high spontaneous fission rate. It has a spontaneous fission rate of 415,000 fission/sec-kg, but emits about 1,000,000 neutrons/sec-kg since each fission produces about 2.2 neutrons. This rate is over 30,000 times that of Pu-239. A contamination of only 1% produces so many neutrons that implosion systems are required to produce efficient bombs. Standard weapon grade plutonium requires a Pu-240 content of no more than 6.5%. Higher levels guarantee that pre-detonation (and reduced yields) will occur even with high performance implosion.

Pu-240 is actually fairly fissile, somewhat more so than U-235. Nonetheless, high concentrations of Pu-240 raise the required critical mass thus aggravating the neutron background problem. Due to its relatively short half-life (1/4 that of Pu-239) Pu-240 produces a correspondingly higher decay heat output (7.1 watts/kg) thus increasing cooling problems in bomb design. Its specific activity is 0.227 Ci/g.

6.2.2.7 Pu-241

This isotope is about as fissile as Pu-239, has a low neutron emission rate, and a moderate heat production rate, and thus does not adversely affect weapon usability of plutonium directly. It decays with a short 14.1 year half-life though, into americium-241 which is non-fissile and does produce a great deal of heat: 106 W/kg. If a weapon thus does include Pu-241 initially over several years or decades its reactivity will decline somewhat, which must be taken into account when the weapon is designed to avoid reduced yield, and its self-heating will increase. The reason that Pu-241 does not emit much heat (3.4 W/kg) despite its very short half-life is its weak beta emission decay mode. Its specific activity is 106 Ci/g.

6.2.2.8 Pu-242

This isotope has a high neutron emission rate, 8.4x10^5 fissions/sec-kg (twice that of Pu-240), and is non-fissile. A substantial concentration thus has the serious adverse effect of driving up the critical mass, while also adding a high neutron background. It has a long half-life, and has a relatively low neutron capture cross section, so it tends to accumulate in recycled reactor plutonium. Its specific activity is 4.0 milliCi/g.

6.2.2.9 Weapon Grade Plutonium

This term is used by the U.S. for plutonium with a Pu-240 content of less than 7%. Typical assays of weapon grade plutonium are given below. The first two are average assays of weapons grade plutonium produced at Hanford, and Savannah River in June 1968. The third is based on soil samples taken outside the Rocky Flats Plant in the 1970s, and is adjusted for the americium-241 also present (the decay product of Pu-241).

TYPICAL WEAPONS GRADE PLUTONIUM COMPOSITIONS

	Hanford (avg. 6/68)	Savannah (avg. 6/68)	Rocky Flats Soil (avg. 1970s)
Pu-238	less than 0.05%	less than 0.05%	trace
Pu-239	93.17%	92.99%	93.6%
Pu-240	6.28%	6.13%	5.8%
Pu-241	0.54%	0.86%	0.6%
Pu-242	less than 0.05%	less than 0.05%	trace

The U.S. has also produced supergrade plutonium with Pu-240 content of 3%, for use as an enricher for lower grade plutonium, and perhaps as an ingredient in special weapon designs. Some U.S. designs have required plutonium with a Pu-240 content as low as 1.5%.

An important question is what the designation as "weapons grade" actually means. The prevalent interpretation has been that this indicates that plutonium with a Pu-240 content less that 7% is actually required for successful weapon construction, or at least, there is a serious compromise of weapon performance above this level.

The Pu-240 content definitely does have weapon design consequences, since it determines the neutron background level and has secondary effects in increasing critical mass (slightly) and thermal output. The neutron background constrains the design by limiting the amount of plutonium included, and requiring implosion speeds above some specified threshold. As noted above, some U.S. weapons designs (presumably older ones) required low Pu-240 contents for these reasons.

However, it is now apparent that these issues are unimportant in advanced weapons designs used by the U.S. since at least the early 1960s. Recently declassified government documents (WASH-1037 Revised, *An Introduction to Nuclear Weapons*, June 1972) make it clear that the designation "weapons grade" is purely an economic one. The cost of plutonium goes down the higher the Pu-240 content. On the other hand, the critical mass goes up with higher Pu-240 contents. Around the 6-7% Pu-240 level, the total cost of the plutonium in the weapon is at a minimum.

This does not mean that plutonium with higher levels of Pu-240 can be used in EXISTING weapon designs. These have been optimized for the use of a specific material and would probably suffer in performance if a different plutonium composition were used.

Assuming an average composition of 93.4% Pu-239, 6.0% Pu-240, and 0.6% Pu-241 (with negligible amounts of other isotopes), the following properties of WG plutonium can be calculated. The initial heat output of freshly prepared WG-Pu would be 2.2 W/Kg, and the spontaneous fission rate would be 27,100 fissions/sec. This fission rate permits a weapon using 4-5 kg of plutonium to be assembled with a very small probability of predetonation by a good implosion system. Over the course of a couple of decades, most of the Pu-241 will decay into Am-241, eventually raising the heat output to 2.8 W/kg. Since Pu-241 is highly fissile, but Am-241 is not, this reduces the reactivity margin in the weapon slightly and must be taken into account in weapon design.

The neutron emission for a 5 kg of WG-Pu, 3x10^5/sec, represents an exposure of some 0.003 rad/hr at 1 m. This is reduced by the reflector and explosive surrounding it, a light-weight weapon might attenuate it by a factor of 5-10. The high RBE for neutrons on the other hand, enhances the risk. Constant close contact with a weapon during a normal work schedule would result in radiation exposures close to the annual occupational limit. Weapon plant employees who handle plutonium cores directly or in glove boxes have limited shielding and may need to be rotated to other tasks to keep exposures down.

Due to the small mass difference between Pu-239 and Pu-240, it has not been considered practical to strip Pu-240 by the common enrichment schemes used with uranium. This has been done for small quantities of plutonium using electromagnetic separation for research purposes. There is no major reason for a sophisticated nation to reduce the Pu-240 content below 6% since it still allows the construction of efficient, reliable fusion bomb triggers. Very low Pu-240 content would allow some additional flexibility in weapon design, which may be desirable for specialized or exotic designs.

6.2.2.10 Reactor Grade Plutonium

For reasons of economics, power production reactor fuel has very high fuel burn-ups. The plutonium from discharged power reactor fuel has a high level of contaminating isotopes. Compositions vary with reactor design, and the exact operating history, but typical ones are:

Isotope	Light Water Reactors Typical	CANDU Reactor 33000 MWD	MAGNOX Reactor 7500 MWD	3000 MWD
Pu-238	2%	1.5%	low	0.1%
Pu-239	61%	56.2%	66.6%	80.0%
Pu-240	24%	23.6%	26.6%	16.9%
Pu-241	10%	14.3%	5.3%	2.7%

Pu-242	3%	4.9%	1.5%	0.3%

The 33000 MWD burnup figure is based on the 3% enrichment commonly used in the 1970s and 80s. With the current cheapness of enriched uranium (based on a world wide post-arms race glut of enrichment capacity), higher enrichment fuel is now being used (4-4.5%), allowing burnups of 45000 MWD or even higher. This results in higher concentrations of Pu-238, 240, 241, and 242.

Using the light water average as the basis, we can calculate a heat output of 14.5 W/kg rising to 19.6 W/kg in 14 years as the Pu-241 decays, reaching an eventual maximum near 24 watts. The neutron emission rate is 350,000 neutrons/kg. The specific activity is 11.0 Ci/g (0.442 Ci/g alpha activity).

Taking into account the effect of isotopic dilution on critical mass (both Pu-239 and Pu-241 are fissile, the others are less so) a bomb fashioned from 8 kg of this material would put out 116 watts (a light bulb of similar size and power is too hot to hold), and 2.8 million neutrons/sec. Using this material in a bomb would be a challenge. Continual active cooling would be needed to prevent deterioration and damage to the core, explosives, and other components. The high rate of neutron emission means that predetonation is inevitable, even with a very efficient implosion system. However, even the relatively primitive Fat Man design would have produced a 0.5 kt or so yield with this material. With optimal implosion design yields in the range of at least several kilotons are possible. If fusion boosting is used, then the adverse effects properties of reactor grade plutonium can be completely overcome, allowing its use in efficient high-yield designs although the material would be less convenient to use.

While reactor grade plutonium would probably be of no interest to a nation with access to better grade material, it could be effectively used by a nation capable of good weapon design, but without access to better fissile material. Even a low technology nation could fashion powerful weapons from it, after all even a 1 kt device greatly exceeds the destruction of any conventional weapon.

Over long periods of time, several decades or centuries, the heat output of reactor grade fuel diminishes greatly from the decay of Pu-238 and Am-241. This has little effect on the neutron background however. Currently, used reactor fuel is commonly stored indefinitely in containment pools near the reactor site. It will remain a proliferation risk until otherwise disposed of, in fact as the fission product radiation levels and plutonium heat output declines, the proliferation risk increases.

Storage for 40 years allows 30% of the Pu-238 and 88% of the Pu-241 to decay. The plutonium composition would then be: 1.5% Pu-238, 67.3% Pu-239, 26.4% Pu-240, 1.3% Pu-241, and 3.3% Pu-242. This would decrease the heat output to 11.7 W/kg and largely stabilize it from post-fabrication increase (maxing out at 13.8 W/kg). Storage for 150 years would result in a

composition of 0.66% Pu-238, 69.06% Pu-239, 26.86% Pu-240, 0.01% Pu-241, and 3.41% Pu-242; for a stable heat output of 7.5 W/Kg.

Using uranium enrichment technology to strip out undesirable isotopes is a real possibility. It is complicated by the presence of multiple isotopes, each separated by one atomic mass unit, compared to the 3 amu difference in uranium. In general this reduces the separation parameter (s - 1) by a factor of 3, and the separation capacity of a stage or plant by a factor of 9. Also a double enrichment process might be necessary. After separating Pu-240 and higher isotopes, a second enrichment might be needed to strip out Pu-238 (depending on the content of the starting material, and how objectionable the heating effect is). The toxicity, neutron emission, and self-heating of the feed, waste, and product would also complicate enrichment operations, compared to uranium.

On the other hand the amount of feed that must be processed to produce a bomb is more than two orders of magnitude smaller than natural uranium. This is due both to the high Pu-239 content (60-70% compared to 0.72%), and the smaller critical mass required (6 kg vs 15 kg). Even with the complications mentioned above, an enrichment plant for upgrading reactor plutonium would be much smaller than a natural uranium plant regardless of the technology used. This material must thus be considered a serious proliferation risk on the national level.

It should be fairly straightforward to produce weapon grade plutonium from reactor grade by electromagnetic separation. Due to the extremely high enrichment factor of this process only one stage of enrichment is needed, and plant output is proportional to the concentration of desired isotope in the feedstock. An electromagnetic separation facility of capable of producing 0.5 bombs a year (similar to the facility that Iraq was planning before the 1991 war) could produce over 100 bombs a year if reactor grade plutonium were available.

Gas diffusion and centrifuges are also viable candidates. The properties of plutonium hexafluoride are similar to uranium hexafluoride, and requires only small changes in diffuser or centrifuge design. If 60% Pu-239/25% Pu-240 material were used as feed, with a product of 94% Pu-239 and tails of 50% Pu-239 (product contains half of the feed Pu-239), then only 2 kg-SWU would be required to produce 1 kg of weapon grade plutonium. This is 1% of the effort required to produce 1 kg of 90% U-235 from natural uranium. Taking into account to the smaller number of SWUs, and the offsetting smaller critical mass of plutonium, a plutonium enrichment plant could produce 29 times as many bombs as an equivalent U-235 plant. A centrifuge plant might need only few dozen stages for a cascade.

The AVLIS technology offers the possibility of inexpensive separation and could have used commercial reactor grade fuel as a feed. This was perhaps the real motivation for pursuing AVLIS research during the eighties.

6.2.2.11 Denatured Plutonium

If plutonium is extracted from reactor fuel and reused in power reactors, its composition becomes progressively less desirable for weapons use. After several fueling cycles the buildup

of Pu-238, Pu-240, and Pu-242 makes it difficult to use in weapons. It has been suggested that blending such recycled fuel is a useful means to "denature", or render proliferation-resistant, supplies of retired weapons-grade plutonium.

This is mainly a deterrent against the use in low technology designs. The elevated thermal output and radiation levels merely have nuisance value, and are not serious deterrents although they create significant design constraints and handling problems. Advanced designs and adequate handling facilities can completely overcome these obstacles. Even at low technology levels destructive devices could be constructed with this material.

6.2.3 Thorium (Th)

Named after Thor, the Norse god of thunder, element 90 was discovered in 1828 by Swedish chemist and minerologist Baron Jons Jakob Berzelius. Berzelius isolated the element from a black silicate mineral, subsequently named named thorite, from the island of Lövö near Brevig, Norway. Thoriums radioactivity was discovered independently in 1898 by Madame Marie Curie and Gerhard Carl Schmidt.

Thorium is a silvery-white lustrous metal that resists oxidation when pure, but typically tarnishes slowly to black over time. Thorium is slowly attacked by water, but does not dissolve readily in most common acids, except hydrochloric. It is slightly soluble in sulfuric acid and nitric acids. Powdered thorium metal is often pyrophoric and should be carefully handled. When heated in air, thorium turnings or scraps ignite and burn with a brilliant white light. Pure thorium is soft, very ductile, malleable, and can be readily worked (cold rolled, swaged, etc.), but drawing is difficult because of thorium's low tensile strength. The oxide content strongly affects thorium's mechanical properties; even pure thorium samples typically contain a few tenths of a percent of thorium oxide. When strongly heated it reacts with halogens, sulfur, and nitrogen. Its electronic configuration is 2-8-18-32-18-10-2 or (Rn)6d2-7s2.

Its principal physical properties are: Density 11.72

Undergoes cubic->body centered cubic crystal change at 1400 C

Melting point 1750 C

Boiling Point 4790 C

Thorium oxide (known as thoria) has an extremely high melting point, 3300 C, the highest of all oxides and higher than the melting points of all but a few substances. This property once accounted for thorium's major commerical use as a refractory ceramic - principally in ceramic parts, investment molds, and crucibles. Its most commonly encountered use was in the Welsbach mantle for "camping" lanterns which produce a dazzling light when heated to very high temperatures.

Thorium is rather abundant. It is the 39th most abundant of the 78 common elements in the Earths crust, at 7.2 parts per million. It is about three times more abundant than uranium, about

equal to lead. It is associated with uranium in igneous rock. As the primary thorium minerals are more resistant to geochemical and physical weathering, the thorium/uranium ratio in sedimentary rock is typically higher than in igneous source rock. Thorium occurs in several minerals, the most common being monazite (a rare-earth-thorium-phosphate mineral, containing from 3 to 9% ThO_2) and thorite (thorium silicate). It also occurs as orangite, a variety of thorite; and in thorianite, a mineral composed of thorium oxide and uranium.

It is estimated that the potential energy represented by accessible thorium deposits exceeds that of uranium and fossil fuels combined by a wide margin. The major deposits of thorium are usually in the form of monazite sands. Large deposits of thorium are found in several parts of the world including the U.S. (New England and North Carolina) and India, which has the largest reserves.

Most thorium production is as a byproduct of rare earth extraction from monazite, which is mined and concentrated in turn as a byproduct of processing heavy-mineral sands for titanium, zirconium, or tin. Only a small portion of the thorium produced was consumed; most was discarded as radioactive waste, an operation that has become increasingly expensive as safety standards regarding radiation have been tightened. The major monazite-producing countries were India, Brazil, Malaysia, China, and Sri Lanka. In 1997 worldwide monazite concentrate production was 7000 tonnes, with a value of US$0.73/kg (calculated on the basis of rare earth content). In 1997 there was no domestic U.S. production. The market for thorium is very small. U.S. domestic consumption of refined thorium products in 1997, according to the U.S. Geological Survey (USGS), was 13.0 tonnes of thorium oxide equivalent valued at only $300,000.

The commercial use of thorium has been declining due to increasingly strict radiation regulations, causing business and industry to replace thorium with non-radioactive substitutes. Thorium is no longer used in the U.S. to make gas lantern mantles, and its use as a refractory ceramic is also declining.

Important uses remain in the production of thoriated tungsten welding electrodes, which are employed used to join stainless steels, nickel alloys, and other alloys requiring a continuous and stable arc to achieve precision welds.

Due to its ability to emit electrons at relatively low temperatures when heated in a vacuum thorium is used in a variety of electronic tube devices. It is used in the thoriated tungsten negative poles of microwave generating magnetron tubes. It is also useful in certain high speed, high power vacuum tube switches that may have seen use in nuclear weapon firing circuits.

Thorium has been widely used in alloys with magnesium by the aerospace industry. These alloys are lightweight and have high strength and excellent creep resistance at elevated temperatures. At least one U.S. missile nuclear warhead, the W40 (or else perhaps the Bomarc missile that carried the warhead), has used thorium-magnesium alloy for structural reasons. Thorium-free magnesium alloys with similar properties have been developed and are expected

to replace most of the thorium-magnesium alloys presently used. Small quantities of thorium were used in dispersion-hardened alloys for high-strength, high-temperature applications.

Thorium is used in other applications such as chemical catalysts (paints, fuel cell elements), elements in special use light bulbs, thoria-containing high-refractivity glass, photo-conductive films, radiation detectors, and target materials for X-ray tubes.

Several methods are available for producing thorium metal; it can be obtained by reducing thorium oxide with calcium, by electrolysis of anhydrous thorium chloride in a fused mixture of sodium and potassium chlorides, by calcium reduction of thorium tetrachloride mixed with anhydrous zinc chloride, and by reduction of thorium tetrachloride with an alkali metal.

In 1997 thorium oxide prices were quoted at $65.55 per kilogram (base price), $82.50 for 99.9% purity and $107.25 per kilogram for 99.99% purity. Thorium alloy prices were no longer available due to lack of sales.

Government stocks of thorium nitrate in the National Defense Stockpile (NDS) were 3,217 tonnes (as thorium oxide equivalent) on 31 December 1997, 2,944 tonnes of which was considered surplus and marked for disposal. The NDS only shipped 1,724 kilograms (1.7 tonnes) for waste disposal research in 1997, no stocks of thorium nitrate were sold during the year.

Twelve isotopes of thorium are known, but the natural primordially existing element consists entirely of a single isotope - Th-232. This isotope has a half-life of 1.405 x 10^10 years and is an alpha emitter. Its specific activity is 0.109 microCi/g, about one-sixth of natural uranium. Thorium is radioactive enough to expose a photographic plate in a few hours. The thorium decay chain produces radon-220 gas ("thoron") which is an inhalation radiaton hazard.

Five other thorium isotopes occur naturally in small amounts as a result of production by the decay of U-238, U-235, and Th-232. These are Th-227 (half life 18.72 days), Th-228 (1.9116 years), Th-230 (75,380 years), Th-231 (1.063 days) and Th-234 (24.1 days). Since all of these isotopes but Th-228 are produced by the uranium decay chains, the concentration of these isotopes in natural thorium depends on the relative amount of uranium present and the possible effects of chemical leaching of the protactinium or actinium intermediates. For practical purposes, the only isotopes that are present at significant levels in purified thorium are Th-228 and Th-230 since the others have very short half-lifes, and Th-228 decays considerably after several years of storage.

Thorium is of interest for nuclear purposes due to its ability to breed fissile U-233 through the chain:

Th-232 + n -> Th-233

Th-233 -> (22.2 min, beta) -> Pa-233

Pa-233 -> (27.0 day, beta) -> U-233

There has been interest in the possibility of establishing a thorium/U-233 fuel cycle using breeder reactors, especially in uranium-poor but thorium-rich India. As a nuclear explosive U-233 is about as good as plutonium-239, and has been employed in several tests by the U.S. (e.g. Teapot MET).

A complication in the use of U-233 in weapons are side reactions that produce the potent gamma-ray emitting U-232. U-232 is produced by the reaction:

Th-231 -> (25.5 hr, beta) -> Pa-231

Pa-231 + n -> Pa-232

Pa-232 -> (1.31 day, beta) -> U-232

The Th-231 is produced in turn by two reactions:

Th-232 + n -> Th-231 + 2n

and

Th-230 + n -> Th-231

The first (n,2n) reaction only occurs to any significant extent when neutrons with energies in excess of 6 MeV are encountered. Only a small percentage of fission neutrons are this energetic, and if the thorium breeding blanket is kept in a reactor region where it is only exposed to a well moderated neutron flux (i.e essentially no neutrons above the Th-232 fission threshold of 500 KeV) this reaction can be nearly eliminated. The second reaction proceeds very efficiently with thermalized neutrons however, and minimizing U-232 from this source requires choosing thorium that naturally has a low Th-230 concentration.

Thorium-232 also undergoes fast fission and could thus be used in thermonuclear secondary stage tampers, but it has a markedly lower cross section for this than does U-238 (at 14 MeV

0.35 barns vs 1.14 barns, and for fission spectrum neutrons 0.078 barns vs 0.308 barns) and is thus not a serious contender for this use.

A major problem with using thorium breeding to obtain fissile material is that it is not naturally present in economical reactor fuels, unlike U-238. To use thorium breeding, a highly enriched fissile material (U-235, U-233, Pu-239) must be used as a fuel with thorium included in the reactor mostly for its breeding potential (i.e. it contributes little or nothing to power production itself, although burning of U-233 produced in situ may contribute to power production). On the other hand thermal breeder reactors are possible using a U-233/thorium breeding cycle, especially if heavy water is used as a moderator.

U-232 contamination is not a serious problem for civilian power use, although unshielded exposure to fuel material during fabrication and handling prior to irradiation would have to be limited. The ability of high levels of U-232 to make diversion difficult and undesirable has been promoted as an anti-proliferation feature of the thorium/U-233 fuel cycle.

Thorium is not currently being used in any foreign or domestic commercial reactors. The fuel is not expected to grow due to the current availability of low-cost uranium. India is interested in developing an indigenous thorium breeding fuel cycle due to its vast thorium deposits, and is accumulating plutonium as the start-up fuel for future thorium cycle breeders. It is believed that India has included U-233 breeding as part of its weapons related activities also.

6.2.4 Other Fissile Elements

U-233, U-235 and the various isotopes of plutonium are not the only known elements and isotopes with the necessary nuclear properties to be usable in fission bombs. Several other candidates are described here, of varying degrees of practicality. None of them have properties that are sufficiently attractive to merit serious consideration in favor of the other "big three", but all of the isotopes discussed below have attracted some attention for a variety of reasons.

Few of these isotopes have well measured critical mass values available in the open literature. Most of the critical mass data are estimates made from their measured or theoretically predicted nuclear properties. The literature is rather spotty in the availability and consistency of these estimates. Accordingly, I have suplemented the published figures with my own calculations for comparison purposes.

My critical mass estimates are for bare spheres at the densest STP phase. The calculations were performed using diffusion theory and a one-group representation of neutronic properties. The one-group parameters are fission spectrum averages calculated from the authoritative ENDF-6 evaluated nuclear data base. Where available I also list estimates from other sources which are usually based on more sophisticated numerical computations than what I have used.

The one-group calculation method consistently underestimates the true critical mass - primarily because it does not take into account the effects of inelastic scattering in softening the neutron spectrum. The one-group calculated critical mass estimates are thus lower bounds on the true value. Comparison between the one-group calculations and the actual values for the highly

fissile isotopes for which good experimental data is available (U-233, U-235, Pu-239, and Pu-241) shows a consistent underestimate of 70-75% of the true value. For less fissile isotopes, where critical mass estimates have been offered by others (these are mostly calculated estimates also, but with more sophisticated models), the underestimates are more severe (at worst 22-29% of the 'true' value for Pu-242). This too is to be expected because the effects of inelastic scattering is relatively greater in less fissile materials. On the other hand, the estimates for extremely fissile transuranics like californium isotopes should be quite good.

The list below is not exhaustive. In particular it is likely that other fissile isotopes of transuranic elements exist (curium, berkelium, etc.).

6.2.4.1 Protactinium

Protactinium was identified by Kasimir Fajans and O.H. Gohring in 1913 from the short-lived Pa-234m isotope and was named brevium ("brief") due to its short lifetime (1.175 minutes). When the much longer-lived isotope Pa-231 was identified by Austrian-Swedish physicist Lise Meitner and the German physical chemist Otto Hahn in 1918 it was renamed protoactinium (i.e. predecessor of actinium). Pa-231 was independently rediscovered in 1918 by Frederick Soddy, John Cranston, and Sir Alexander Fleck; and by Fajans. The name protoactinium was shortened to protactinium in 1949.

The metal itself was not isolated until 1934 when Aristid V. Grosse developed two methods. One involved reduction of the pentoxide Pa2O5 with a stream of electrons in a vacuum and the second involved dissociating the iodide PaI_5 by a heated filament under a high vacuum.

Protactinium is a malleable silvery metal with a bright metallic lustre which it retains for some time in air, tarnishing with an oxide coating. Its has a tetragonal crystal structure and a density is 15.37. Its melting point is 1570 C, its boiling point is 4227 C. Since its half-life is similar to Pu-239, it has similar alpha activity as well and must handled with similar safety precautions.

Protactinium has three valence states: III through V. Protactinium in most of its compounds exhibits an oxidation state of V (thus resembling tantalum). It is soluble in dilute HF. Its compounds readily hydrolyze in water, forming colloids, but dissolve by forming complex ions (as with the fluoride ion in hydrofluoric acid). It forms oxides, and halides, and reacts with H2 at 250-300 C to form PaH3. Its electronic configuration is 2-8-18-32-20-9-2 or (Rn)5f2-6d1-7s2.

Protactinium is one of the rarest and most expensive naturally occurring elements. Protactinium has 20 isotopes, its predominant and stablest isotope, protactinium-231, is a decay product of U-235 and exists in decay equilibrium with it in uranium ores. The creation and destruction is governed by the reaction chain:

U-235 -> (7.1 x 10^8 yr, alpha) -> Th-231

Th-231 -> (25.52 hr, beta) -> Pa-231

Pa-231 -> (32,500 yr, alpha) -> Ac-227

From the half-life ratio of U-235 and Pa-231 we can calculate that they exist in the ratio of (3.25 x 10^4)/(7.1 x 10^8) = 4.6 x 10^-5. Since 0.72% of natural uranium is U-235, the ratio for Pa-231 to all uranium atoms is 3.3 x 10-7 (3.2 x 10-7 by mass), discounting any chemical separation of the elements. Typical pitchblende ores contain about 0.1 ppm of Pa-231. Some ores from Congo (formerly Zaire) have up to 3 ppm.

Based on theoretical considerations Pa-231 in the past has been considered a potential weapon-usable material, the only such material with an atomic number lower than that of uranium. It was the only candidate fissile material that could in principle be obtained from natural sources through conventional chemical separation techniques alone. It was considered as a candidate fissile material during the early stages of atomic research in the 1940s.

This possibility has now been effectively excluded. State-of-the-art calculations by IRSN (L'Institut de radioprotection et de sűreté nucléaire) have shown that Pa-231 is probably not critical even in an infinite metal assembly (Report C4/TMR2001/200-1).

Pa-231 is extracted from uranium ores and can be concentrated from uranium ores as part of the normal process of ore refining. If Canada, the world's largest uranium producer, had done this with its 1994 uranium production (9700 tonnes) it would have obtained 3.1 kg of Pa-231. Accumulating an unreflected critical mass in this manner would take several decades.

Unless the "tails" piles left from uranium mining and extraction were stringently guarded (which would itself be suspicious) it would be impossible to conceal such activity since an examination of tails samples would immediately reveal if Pa-231 accumulation was being practiced.

Perhaps motivated by interest in its possible use in weapons, between 1959 and 1961 the Great Britain Atomic Energy Authority extracted by a 12-stage process 125 g of 99.9% protactinium, the world's only stock of the metal at the time. The extraction was made from 60 tons of waste material at a cost of about $500,000.

Protactinium can also be produced from natural thorium in a reactor by the following reactions:

Th-232 + n -> Th-231 + 2n

Th-231 -> (25.52 hr, beta) -> Pa-231

6.2.4.2 Neptunium

Neptunium was the first transuranium element discovered, and the first synthetic transuranium element to be prepared. Though traces of neptunium have subsequently been found in nature, it was discovered by Edwin M. McMillan and Philip H. Abelson in May 1940 at Berkeley, California, USA, who bombarded uranium with neutrons produced from a cyclotron, producing

Np-239. Since it was the next element in the periodic table after uranium, by analogy it was named after the planet Neptune, which is the next planet out from the Sun after Uranus (this pattern was subsequently followed in the naming of plutonium).

As is also true of plutonium, trace quantities of the element are actually found in nature due to neutron-induced transmutation reactions in uranium ores produced by spontaneous fission neutrons.

Neptunium is a silvery metal, with a density of 20.45 (Np-237, 25 C), neptunium is the fifth densest element. It has at least three allotropic (crystalline) forms -- the alpha form which is the densest (stable below 280 C, orthorhombic, density 20.45 at 25 C), beta (above 280C, tetragonal, density 19.36 at 313 C), and gamma (above 577 C, cubic, density 18.0 at 600 C). Its melting point is 637 C, and its boiling point is estimated at 5235 C.

Neptunium is chemically reactive and similar to uranium with oxidation states from +3 (III) to +7 (VII). The metal is prepared by the reduction of NpF3 with barium or lithium vapor at about 1200C. Neptunium of the five ionic oxidation states the pentavalent is the most stable in solution. Neptunium ions in aqueous solution possess characteristic colours: Np3+, pale purple; Np4+, pale yellow-green; NpO+2, green-blue; NpO2+2, varying from colourless to pink or yellow-green, depending on the anion present. The element forms tri- and tetrahalides such as NpF3, NpF4, NpCl4, NpBr3, NpI3, and oxides for the various compositions such as are found in the uranium-oxygen system, including Np3O8 and NpO2. Its electronic configuration is 2-8-18-32-22-9-2 or (Rn)5f4-6d1-7s2.

Fifteen isotopes of neptunium are now recognized. The isotope neptunium-237 (discovered in 1942 by Glenn T. Seaborg and Arthur C. Wahl) is the longest lived, with a half-life of 2.144 x 10^6 years, and is fissile. It is most easily produced by the following reactions:

U-235 + n -> U-236

U-236 + n -> U-237

U-237 -> (6.8 days, beta) -> Np-237

It is also produced by the decay of Am-241, but the 433 year half-life of this latter isotope makes the production rate very small.

Published estimates for the critica mass of neptunium-237 are around 90 kg (range of estimates is 75-105 kg), my one-group lower bound calculation is 19 kg. It has a very low spontaneous fission rate (< 5x10^-2 F/sec-kg). The high critical mass value (almost double that of HEU) and

the high cost of manufacture makes it unattractive for weapons use. A certain amount Np-237 is produced normally by U-235 capture in reactors. A typical power reactor can discharge about 0.4 kg of Np-237 per ton of fuel. Fast reactors can discharge a significantly higher amount.

Np-237 is used in some types of neutron detection instruments

6.2.4.3 Americium

Americium (atomic number 95) was the fourth transuranic element to be synthesized (curium, atomic number 96, was discovered a few months previously). It was identified by Glenn T. Seaborg, Ralph A. James, Leon O. Morgan, and Albert Ghiorso in 1944 as the result of successive neutron capture reactions by plutonium isotopes in a nuclear reactor to form Am-241. Americium was also prepared by Glenn Seaborg, et al by bombarding Pu-234 with alpha particles to produce Am-241.

It is a silvery metal, the luster of freshly prepared americium metal is whiter and more silvery than plutonium or neptunium prepared in the same manner. It appears to be more malleable than uranium or neptunium. Americium tarnishes slowly in dry air at room temperature. It has two allotropic forms. The alpha form is stable at low temperatures and is a double close-packed hexagonal structure, density 13.67 (molar volume 17.6 cm^3/mole, 20 C). It transitions at 1074 C to the beta form with a face-centered cubit structure. The melting point is 1175 C.

Americium has valences III, IV, V, and VI, in acidic aqueous solution with the following ionic species: Am3+, pink; Am4+, rose (very unstable); AmO2+ , yellow; and ((AmO2)+2)2, light tan. Trivalent americium is most common in aqueous solution and in this state is very similar to the other actinide and lanthanide elements. There is some evidence that the ion Am2+ has been prepared in trace amounts, its existence suggesting that americium is similar to its lanthanide homologue, europium, which can be reduced to its divalent state. Am IV is known only in the solid state. Americium reacts with oxygen to form the dioxide AmO2 and with hydrogen to form the hydride AmH2. Americium dioxide is obtained by ignition of of most trivalent compounds, AmF_4 is produced by fluoridation of the dioxide or trifluoride. Its electronic configuration is 2-8-18-32-25-8-2 or (Rn)5f7-7s2.

Am-241 has a half-life of 433 years, and an estimated bare critical mass of 84 kg, my one-group lower bound calculation is 23 kg. It has a moderate spontaneous fission rate (2.4x10^2 F/sec-kg).

The most commonly used isotope of americium is Am-241, now produced in quantity by the decay of Pu-241:

Pu-241 -> (13.2 yr, beta) -> Am-241

Since Pu-241 is normally present in freshly made weapons-grade plutonium, Am-241 accumulates in the material as the Pu-241 decays.

Since Pu-241 is normally present in freshly made weapons-grade plutonium, Am-241 plays a significant role in weapons issues. Freshly made "weapons-grade" plutonium normally contains 0.5 - 1.0% Pu-241, reactor grade plutonium typically contains 5-15% Pu-241 to as much as 25%. In the course of a few decades, most of this Pu-241 will decay into Am-241. Am-241's energetic alpha decay and relatively short half-life give it a high specific activity and heat output. Much of the alpha and gamma activity and heat produced in older weapon grade plutonium is due to Am-241. Reprocessing WG-Pu reclaimed from old weapons eliminates this problem.

This high critical mass (almost double that of HEU) and relatively high cost make it unattractive for use as a weapon fissile material. It is made even more unattractive by the fact that a bare critical mass of Am-241 produces 9.6 kw of heat! It would be challenging to make a storable weapon that can tolerate this kind of heat output, probably requiring continuous active cooling in storage.

The longer lived isotope Am-243 (half-life 7370 years) is another possible fissile material. It has a higher critical mass, published estimates are unavailable but my one-group calculation is 38 kg. By comparison with the estimates for Am-241 this suggests a true value in the vicinity of 140 kg, which would have a heat output of 900 watts. Am-243 is a much less available and more expensive material however since it is produced by neutron irradiation of Am-241.

The isotope americium-241 is the most important because of its availability. It has been prepared in kilogram amounts from plutonium and has been used industrially in fluid-density gauges, thickness gauges, aircraft fuel gauges, and distance-sensing devices, all of which utilize its gamma radiation. Am-241 is used in very small amounts in smoke detectors (0.26 micrograms per detector) where its alpha emissions are employed. It has also been used as a diagnostic aid in bone mineral analysis. In Sept. 1998 the ORNL Isotopes Division charged $728.00 per gram of americium-241, and $180.25 per milligram of americium-243.

6.2.4.4 Californium

Californium, the sixth transuranium element to be discovered, was produced in 1950 by Glenn T. Seaborg, Stanley G. Thompson, Albert Ghiorso, Kenneth Street in 1950 by bombarding microgram quantities of Cm-242 with 35 MeV helium ions in the Berkeley 60-inch cyclotron to produce Cf-245. It was named after the State and University of California.

Metallic californium (Cf-249) has been prepared by reducing Cf_2O_3 with lanthanum metal. Two crystalline forms have been identified, a face-centered cubic and hexagonal close-packed structures. The melting point is 900 C (+/- 30). From theoretical models the density is 15.3 for the Cf-252 isotope.

Californium has valences II, III, and (rarely) IV. Californium (III) is the only ion stable in aqueous solutions, all attempts to reduce or oxidize californium (III) in solution having failed. In addition to the Cf(III) oxide (Cf2O3), several binary compounds with halogens are known -- the Cf (II) bromide and iodide; the Cf (III) fluoride, chloride, bromide and iodide; and the Cf(IV) fluoride.

Other compounds that have been prepared include the oxychloride, CfOCl. Its electronic configuration is 2-8-18-32-28-8-2 or (Rn)5f10-7s2.

Known isotopes of californium include:

Isotope	Atomic Mass	Half-life		Decay Mode
Cf-242	242.06372	3.3	m	alpha to Cm-238
Cf-243	243.065	11	m	alpha to Cm-239
Cf-244	244.06599	20	m	alpha to Cm-240
Cf-245	245.06807	44	m	alpha to Cm-241
Cf-246	246.06884	36	h	alpha to Cm-242; SF (0.2%)
Cf-247	247.0712	3.11	h	electron capture to Bk-247 (99.97%)
Cf-248	248.07218	334	d	alpha to Cm-244; SF (0.013%)
Cf-249	249.07485	351	y	alpha to Cm-245; SF (5.2E-7%)
Cf-250	250.07640	13.1	y	alpha to Cm-246; SF (0.079%)
Cf-251	251.079580	898	y	alpha to Cm-247
Cf-252	252.08162	2.64	y	alpha to Cm-248; SF (3.08%)
Cf-253	253.08513	17.8	d	alpha to Cm-249; beta to Es-253
Cf-254	254.08732	60.5	d	SF (above 99%); alpha to Cm-250 (less than 1%)
Cf-255	255.0910	1.4	h	beta to Es-255

The best known and most widely used isotope of californium is Cf-252. This isotope has an extremely high rate of spontaneous fission, a remarkable 3.08% of decay events being fissions, and a relatively short half-life (2.64 years) making it a potent emitter of fission spectrum neutrons (one microgram emits 170 million neutrons per minute). The half-life is long enough that a sample of Cf-252 can remain active for several years, allowing it to serve as a convenient compact portable neutron source. It has found use as a start-up neutron source for nuclear reactors, and in neutron moisture gauges and in well-logging (the determination of water and oil-bearing layers in bore holes). It is also being used as a portable neutron source for discovery of metals such as gold or silver by on-the-spot activation analysis. Cf-252 foils are used as a source of fission fragments for research purposes. Cf-252 has also been used as a source of

fisson neutrons in a mock-up of the Little Boy bomb in Hiroshima dosimetry studies. As of May 1975 some 63 mg of Cf-252 had been produced and sold by ORNL.

Californium has acquired a reputation for having a remarkable small critical mass, usually alleged to be "in the gram range", creating much speculation about possible use in "pocket nukes" -- very small fission weapons. While this reputation for a small critical mass is partly justified, it is also exaggerated and appears to have taken on the character of an urban legend. In particular the assertion that californium has a critical mass "in the gram range" is at best very misleading.

As is true of other transuranic elements, the odd-numbered isotopes of californium have the most desirable fissile properties. Only two isotopes are plausible candidates for fission explosives - the long lived isotopes Cf-249 and Cf-251. Popular speculation about californium as a weapon material has usually centered on Cf-252 due to its neutron-emitting celebrity. Its comparatively short half-life (for weapons purposes), intense and penetrating neutron radioactivity, and high thermal output make it quite unsuitable for this however.

I have seen no actual published estimates for californium isotope critical masses, so I have done one-group calculations for the two californium isotopes that are most plausible for weapons use (Cf-249 and Cf-251) as well as the often discussed Cf-252 (see the beginning of this section for more on this method). I used theoretical estimates for the density of californium based on its position in the periodic table, which are fairly reliable, since experimental data is unavailable. The estimation model used tends to underestimate critical mass sizes but should be quite accurate for highly fissile materials like thse isotopes. The calculated bare sphere critical masses are:

- Cf-249 5.90 kilograms
- Cf-251 1.94 kilograms
- Cf-252 2.73 kilograms

By using a thick beryllium reflector the critical masses can be reduced to 40% or so of their bare value so that the most fissile of these isotopes, Cf-251, would have a reflected critical mass of 780 grams. This is arguably "in the gram range" since it is less than one kilogram, but it is obviously much, much closer to one kilogram than to one gram. Implosive compression can reduce this further. Just as a powerful and heavy implosion systems can produce low yield nuclear explosions from as little as 1 kg of Pu-239 (yield up to 100 tons, with the implosion system weighing on the order of 1000 kg), a relatively large implosion system could produce a low yield explosion from as little as 200 grams of Cf-251. The yield would be proportionately smaller of course, around 20 tons, and the total mass on the order of 200 kg. The U.S. has tested devices with 20 ton yields using the vastly cheaper plutonium (though much more of it) which only weighed 13 kg.

Experience with small nuclear devices indicates that it is impossible to make a nuclear device with a total mass less than the bare sphere critical mass of the fissile material used. Beryllium reflectors and high explosives can reduce the fissile mass required as indicated, but at the

expense of adding more weight than is saved. Thus a nuclear device smaller than 2 kilograms or so using Cf-251 is almost certainly impossible.

The isotope Cf-249 results from the beta decay of Bk-249 while the heavier isotopes are produced by intense neutron irradiation of this product. Logically one would expect that Cf-252 would be more expensive then than Cf-249. However, in Sept. 1998 the ORNL Isotopes Division charged $180.25 per microgram of californium-249, and $56.00 per microgram of californium-252, indicating that factors other than ease of prodcution were setting the price. If Cf-251 could be manufactured in quantity for the 1998 sale price of Cf-252, then a critical mass would cost $11 billion.

6.3 Fusionable Materials

Fusion fuels are all light elements, which are enormously more abundant than the heavy elements relevant to fission. Isotopic separation is fairly easy in light elements since low masses mean relatively large mass ratios (Li-7 is only one atomic mass unit heavier than Li-6, but it is 17% heavier; U-238 is 3 amu heavier than U-235, but it is only 1.3% heavier).

6.3.1 Hydrogen Isotopes

Hydrogen was identified as a distinct substance by Henry Cavendish in 1766, and was named by Antoine Laurent Lavoisier. Natural hydrogen consists of two isotopes: H-1 (light hydrogen or protium), and deuterium (D, H-2, heavy hydrogen) which occurs in nature in a concentration of 0.015% (one atom in 6760 of light hydrogen). Light hydrogen participates in fusion reactions extremely slowly (that's why the sun is still around). Deuterium fuses much more readily. In the smallest stars, known as brown dwarfs, only deuterium fusion can occur and once this is exhausted the star becomes inert and planet-like. All of the deuterium in the universe today was created in the first three minutes of the Big Bang, it has been slowly depleted by stellar burning. The unstable super heavy isotope tritium rapidly decays and thus exists in nature only in minute quantities.

In elemental form all hydrogen isotopes are gases with very low densities and boiling points. This often makes them inconvenient to incorporate into practical weapons, so lithium deuteride/tritide compounds are frequently used. *6.3.1.1 Deuterium (D)*

This natural isotope was discovered by American chemist Harold C. Urey and his associates F.G. Brickwedde and G.M. Murphy in 1931. Urey was awarded the Nobel Prize for Chemistry for this achievement in 1934. Urey had predicted a vapor pressure difference between the molecular hydrogen (H2) and of a corresponding molecule with one hydrogen atom replaced by deuterium (HD) and, thus, the possibility of separating these substances by distillation of liquid hydrogen. The deuterium was detected (by its atomic spectrum) in the residue of a distillation of liquid hydrogen.

Deuterium was first prepared in pure form in 1933 by G.N. Lewis, using the electrolytic method of concentration discovered by E.W. Washburn. When water is electrolyzed--i.e., decomposed by an electric current (actually a water solution of an electrolyte, usually sodium hydroxide, is used)--the hydrogen gas produced contains a smaller fraction of deuterium than the remaining water, and, hence, deuterium is concentrated in the water. Very nearly pure deuterium oxide is obtained when the amount of water has been reduced to about one hundred-thousandth of its original volume by continued electrolysis. This was the standard method of preparation of D2O before World War II.

Deuterium is twice as heavy as hydrogen-1, a fact that is very noticeable in the elemental state. Deuterium oxide, D2O, is commonly called "heavy water".

Comparison Between Natural Hydrogen and Deuterium

	H	D	H2O	D2O
density (sol.)	0.0763(13 K)		0.917 (0 C)	
density (liq.)	0.0700(20.4 K)	0.169(20.4 K)	1.000(4.0 C)	1.106(11.2 C)
dens (gas, g/l)	0.08987(0 C)	0.17960(0 C)	0.5974(100 C)	
melting point	13.96 K	18.73 K	0 C	3.79 C
boiling point	20.39 K	23.57 K	100 C	101.41 C
critical temp.	33.26 K	38.41 K	374.2 C	371.5 C
crit. press.	12.8 atm	16.432 atm	218 atm	

Two interesting properties of heavy water are that it releases significant amounts of heat upon mixing with light water, and that heavy water ice is slightly denser than light water. Because of this latter property, it will sink if dropped into a glass of regular water, making it useful as a (somewhat expensive) scientific puzzle demonstration. Heavy water is toxic only when enough has been consumed to replace a substantial percentage of light water from the body (several liters or more) so one glass with some heavy water ice cubes could be consumed safetly.

Deuterium is primarily manufactured for use as a nuclear reactor moderator. It is the best moderating material known. For reactor use it is enriched to 99.75% purity or higher. The U.S. has not manufactured heavy water since 1982. Its cumulative production is estimated at 7300 tonnes. The U.S. produced heavy water at Dana, Indiana (discontinued in 1957) and at the Savannah River Facility in South Carolina (discontinued 1982). Reactor grade heavy water costs about $500/kg.

Enrichment Methods

The normal procedure is for a heavy water plant to use a combination of separation techniques in series to achieve maximum economy. Dual temperature chemical separation is the method of choice for early stages of enrichment when the concentration is very low. This is because huge volumes of water are handled, and techniques that have very low energy consumption per unit of feedstock are essential. More energy intensive, but faster, methods may be used to achieve high purity.

Chemical Separation of Deuterium

The most economical method uses the fact that D concentrates to varying (small) degrees in hydrogen sulfide (H2S) gas at different temperatures. By sending liquid water and H2S mixtures through alternating hot and cold towers, D progressively builds up, alternating between the gas and liquid phases. This method, called the Girdler-Sulfide (GS) process, is usually employed until a concentration of 15% or so (13.5% at Savannah River) is reached. Since the 1950s the dual temperature H2S process has been the only process known to be used on a large scale anywhere in the world, except for the final concentration stages.

The first process used for deuterium production in association with the U.S. weapons program was fractional distillation of water. The world's first heavy water reactor was loaded with D20 produced by this process. Water distillation was used as the second and final enrichment stage of the Savannah River Plant before shutdown, bringing the concentration to 99.8%. It is widely used around the world as a final stage, and has been probably the most widely used process after the GS process.

During the 1950s the U.S. used and liquid hydrogen distillation to produce deuterium for the hydrogen bomb program since it could be scaled up for production rapidly and has a high separation factor.

Electrolysis has been popular for the final stages of enrichment since it has a very high separation factor (D remains behind as water is electrolyzed), and when the volumes of liquid are small, the high energy cost is unimportant. Deuterium can be concentrated rapidly to 99.8% purity. The U.S. once used it as the final stage at Savannah River, taking 90% enriched feed from the water distillation cascade.

Two deuterium enrichment techniques, electrolysis and liquid hydrogen distillation, can produce substantial quantities of deuterium as a byproduct of industrial chemical processes or other endeavors. The famous Norsk Hydro plant at Vemork, Norway which was sabotaged by the allies during World War II, was an example of an industrial electrolysis plant that produced heavy water as a byproduct.

6.3.1.2 Tritium (T)

Discovered in 1934 by bombarding deuterium with accelerated deuterons (thus causing a fusion reaction), tritium (T, H-3) is radioactive with a half-life of 12.355 (+/- 0.010) yr. It exists in minute amounts naturally but significant quantities must be made artificially.

It is usually produced by the Li-6 + n -> T + He-4 reaction in nuclear reactors designed for isotope production. The lithium is enriched in Li-6 prior to irradiation, and is usually loaded in the form of aluminum canned lithium fluoride slugs or aluminum-lithium alloy rods. The latter type of target has been used at Savannah River. Tritium was been purified by heating the rods in a vacuum to release hydrogen and helium gases, followed by purification by palladium diffusion (to produce pure hydrogen) and cryogenic distillation (to produce pure tritium from the hydrogen isotopic mixture). These same refinement processes has been used in reclaiming tritium from recycled weapons reservoirs. Currently this is done using a new hydride process.

Substantial quantities are also generated in nuclear reactors moderated by heavy water through deuterium capture: D + n -> T. The Canadian CANDU (Canadian Deuterium) commercial power reactors collectively produce 250-500g of tritium annually.

Since it currently has no operating tritium production plant, the U.S. is exploring the possibility of producing tritium using a high energy (approximately 1 GeV) proton accelerator to be built at Savannah River (a scheme known as APT - Accelerator Production of Tritium). The APT approach uses the proton beam to generate neutrons through spallation from heavy nuclei - probably lead or tungsten. Each proton-nucleus collision produces about 30 neutrons, and leaves a single radioactive nucleus behind. At present the favored design for tritium production surrounds the heavy element target with a moderating blanket of heavy water, through which tubes circulate helium-3 gas. He-3 (itself produced by tritium decay) has a very high cross section for thermal neutrons (5000 barns) and forms tritium upon neutron capture. This system in effect recycles the existing tritium inventory. As an alternate approach a matrix of lead and lithium-6/aluminum alloy rods can be used for the target. Planned capacity is 3 kg/tritium a year, upgradable to 5 kg/yr.

T emits very weak radiation, 18.6 KeV betas with no gamma emissions. These are easy halted by thin layers of material (plastic, skin, etc.), which is why tritium is now used for glow-in-the-dark watches instead of radium. Tritium can only be a health hazard if taken internally. Tritium's specific activity is 9649 Ci/g. The heat output from tritium decay is 0.324 W/g.

When in the form of elemental hydrogen, inhaled tritium gas is not hazardous since the body does not readily absorb hydrogen, only about 0.004% of inhaled tritium gas is absorbed. Tritiated water (HTO) in liquid or vapor form is quite hazardous however, when inhaled nearly all of it is absorbed. There is also significant absorption of HTO vapor through the skin. The estimated LD50 (quantity that will kill half of the people absorbing it) is 10 Ci, given the high specific activity this makes the lethal dose about 1 mg. This is comparable to the toxicity of the nerve gas sarin.

6.3.2 Lithium (Li)

Lithium, element 3, is named after the Greek word for stone - lithos. Lithium is one of only three elements to have existed since the beginning of the Universe, having been synthesized during the Big Bang (the other two are hydrogen and helium). It is the lightest metal, both by density and by atomic weight. It was discovered in 1817 by Johan August Arfvedson in the mineral petalite. The first isolation of elemental lithium was achieved later by W.T. Brande and Sir Humphrey Davy by the electrolysis of lithium oxide. In 1855, Bunsen and Mattiessen isolated larger quantities of the metal by electrolysis of lithium chloride. In 1923 the first commercial production of lithium metal was achieved by Metallgesellschaft AG in Germany using the electrolysis of a molten mixture of lithium chloride and potassium chloride.

It is distributed widely throughout nature. Lithium ranks 35th in order of abundance of the elements in the crust of the earth, being slightly less abundant than copper, but more abundant than lead. It is abundant and is mined commercially in many places including North America, and has wide commercial application. Natural lithium consists of two isotopes Li-6 (7.42%) and Li-7 (92.58%).

Lithium is an alkaline earth metal, and is the lightest of all metals. It is is softer than lead but harder than the other alkali metals. It can easily be drawn into wire and rolled into sheets. It is chemically very reactive, igniting readily in air and burning with a brilliant white flame. A freshly cut lithium chunk is silvery white but it tarnishes in a minute or so in air to give a grey surface. Since it corrodes very rapidly upon exposure to air, when it is stored it must be immersed in a non-hygroscopic liquid such as naphtha. Chemically, lithium resembles sodium in its behavior. It has the highest specific heat of any solid element. It has a body-centered cubic crystal structure. Its electronic configuration is 2-1 or 1s2-2s1.

Principal physical characteristics:

Density 0.534

Melting Point 180.54 degrees C

Boiling Point 1342 degrees C

Lithium is found in economically exploitable quantities in such minerals as spodumene, lepidolite, amblygonite, and petalite. After decades as the worlds leading producer of lithium and its compounds, the United States was surpassed in 1997 when Chile became the worlds largest lithium carbonate producer. Lithium has an enormous variety of uses in the chemical industry, it is widely used in small amounts by many other industries. The aluminum, ceramics and glass, lubricating grease, and synthetic rubber industries consumed most of the lithium minerals and chemicals sold in 1997. The largest use of lithium compounds in the United States are lithium carbonate additives in ceramics and glass manufacturing processes. The domestic manufacture of thermal shock-resistant cookware (pyroceramics) consumes the majority of lithium used in the ceramics and glass industry. The manufacture of black-and-white television picture tubes consumes significant amounts of lithium concentrates overseas. In the aluminum industry lithium is used in the electrochemical production process. The third largest end use for lithium

compounds is as catalysts in the production of synthetic rubbers, plastics, and pharmaceuticals. Lithium has recently found extensive use in high energy lithium-ion and lithium-polymer batteries. Lithium carbonate, a common mineral, is used in the treatment of bipolar disorder and some forms of depression.

Lithium metal is obtained by the electrolysis of a mixture of fused lithium and potassium chloride. The metal is useful as an alloying agent with magnesium and aluminum to produce very light alloys. The alloys, which are 2% to 3% lithium by weight, are attractive to the aircraft industries because of their reduced density and superior corrosion resistance compared with those of conventional aluminum alloys. The metal is used as a deoxidizer and to remove unwanted gases during the manufacture of nonferrous castings. Lithium vapor is used to prevent carbon dioxide and oxygen from forming scale in furnaces in heat-treating steel.

Lithium is used in nuclear weapons in the form of lithium hydride. It can be used in weapons simply as a convenient means of storing deuterium fusion fuel (lithium hydride contains more hydrogen per unit volume than liquid hydrogen does), or it can serve as an essential fusion fuel in its own right. Lithium hydrides have also been used in weapons for neutron shielding. For weapons use it is usually desirable to enrich lithium in Li-6 to provide a more effective fusion fuel. This has been done using electromagnetic separation in the U.S. and Russia, but large scale production uses chemical enrichment methods.

The Y-12 Plant at Oak Ridge has been the major enrichment site in the U.S., using an electrochemical exchange process. This process relies on electromotive exchange, a differential affinity of mercury for Li-6. A lithium-mercury amalgam is agitated with an aqueous lithium hydroxide solution. Li-6 concentrates in the lithium-mercury amalgam phase, Li-7 in the hydroxide solution. Counterflow through a cascade of stages produces highly enriched Li-6 in the amalgam. Chemical reaction between the lithium amalgam and the water is prevented by placing anodes in the hydroxide solution, and using the amalgam as a cathode. The depleted lithium is removed from the aqueous solution through electrolysis. The enriched lithium is recovered from the amalgam by decomposing it with water in the presence of a graphite catalyst, producing an aqueous lithium hydroxide solution and regenerated mercury ready for amalgamation with fresh lithium feed material.

Originally the exchange was done using large shallow pans in the Elex (Electro-Exchange) Plant at Y-12, which operated from 1953-1956. The Elex plant began construction in autumn 1952 for use in weapons tests conducted in 1954. The plant reached initial operating capacity on 14 August 1953, with production expanding rapidly a few months later. The Elex process required repeated mixing of the amalgam and aqueous solution in the pans, followed by settling, and was thus not truly continuous.

The vast majority of enriched lithium was prepared at the Y-12 Colex (Column Exchange) plant which operated from 1955 through 1961. This plant substituted packed columns for the pans, and operated by continuous counter-current flow. Enrichment operations ceased entirely in 1963 when a large stockpile had accumulated.

Three enrichment levels were produced: 95.5% Li-6, 60% Li-6, and 40% Li-6. The depleted lithium contained 1-4% Li-6. Lithium enriched to 95% Li-6 was used in the Castle Union test on 25 April 1954 that produced a 6.9 Mt yield (1.9 Mt from fusion). 40% Li-6 was used in the 15 Mt Castle Bravo test on 28 February 1954 (5 Mt fusion). But natural lithium was used in Castle Romeo (26 March 1954), giving a 11 Mt yield (4 Mt fusion).

The U.S. produced 442.2 tonnes of enriched lithium (of various grades) between 1954 and 1963, with 30,917 tonnes of depleted lithium hydroxide monohydrate as tails (this is a lithium content of 5152 tonnes). Much of this depleted lithium has been offered for sale as surplus since 1968. 10,467 tonnes of natural LiOH.H2O remain in inventory. In Sept. 1997 ORNL was offering Li-6 of 95-96% purity for sale at $1.30 per gram. High purity (>99%) Li-7, which is not the normal byproduct of Li-6 enrichment, was offered for $6.70 per gram.

Lithium hydrides are white crystalline solids (m.p. 620 deg C), of generally salt-like physical character. The term "lithium hydride" may refer specifically to a compound of light hydrogen or may be used generically to refer any lithium-hydrogen compound regardless of hydrogen isotope. They are usually prepared by direct reaction between hydrogen and metallic lithium at elevated temperature. Lithium hydrides have low density (0.82 for LiH, 0.92 for LiD). Lithium hydrides have no known solvent. LiH reacts violently with water, producing hydrogen gas (a fact that makes it useful for special purpose hydrogen gas generators, such as those used to inflate lifeboats). Pure lithium hydride quickly turns grey on exposure to light or air due to the formation of hydroxides and carbonates.

It can be formed into ceramic-like pieces for use in weapons, but extreme care is required to keep it dry to avoid deterioration. It is fabricated and handled in special "dry rooms" at Y-12, with workers wearing sealed suits to prevent body moisture from escaping. Fabrication of LiH weapon components has been performed by hot isostatic pressing. The LiH powder is first packed into a sealed plastic jacket. The jacketed LiH is then placed in an oil filled pressurized autoclave where pressures of 10,000 psi or greater and high temperatures compact it into a solid mass. This mass is later machined to its final shape. Thin film hydriding techniques to form thin lithium hydride layers may be used in some modern weapons.

In addition to its use as a thermonuclear fuel, lithium-6 has been used to harden nuclear weapons against outside neutron fluxes (from other nuclear weapons for example). Lithium-6 hydride has been used for this, the light hydrogen isotope presumably for its very strong moderating effect, and Li-6 for its low mass and high neutron absorption cross-section in the 0.1-1.0 MeV range.

6.4 Other Materials

Some other materials of special interest for nuclear weapons use are listed below.

6.4.1 Beryllium (Be)

Beryllium, element 4, is named after the mineral beryl, which is a natural ore (beryllium aluminum silicate, $Be_3Al_2(SiO_3)_6$, aquamarine and emerald are precious forms of this mineral). It was discovered discovered as an oxide by Louis Nicolas Vauquelin in 1798, and was isolated as the metal by Friederich Wöhler in 1828 (and independently by Antonine Alexandre Brutus Bussy) by the action of potassium on $BeCl_2$ in a platinum crucible. Commercial production began in the 1950s as a result of the nuclear weapons program in the U.S.

Beryllium is an alkaline earth metal with a valence of two. It is steel gray and tarnishes only slightly in air, becoming covered with a thin layer of oxide. It is one of the lightest metals, being lighter than aluminum and only slightly heavier than magnesium. Nonetheless is still surprisingly dense considering its very low atomic weight. This is due to its having one of the densest packing of atoms per unit volume of any element (only the diamond form of carbon is higher). All of the beryllium in existence is thought to have arisen from cosmic ray collisions in interstellar space.

Beryllium and its compounds (particularly its oxide) have unique and remarkable combinations of physical, chemical, electronic, and nuclear properties. Accordingly despite its high cost it is widely used in a enormous variety of applications.

Beryllium has only one natural isotope with an atomic weight of 9. Its electronic configuration is 2-2 or 1s2-2s2.

Beryllium's principal properties are:

Density 1.848 (20 C)

Melting point 1278 C (+/- 5 deg.)

Boiling point 2970 C (at 5mm pressure), above 3000 C at standard atmospheric

Beryllium has the unusual property of that many of its salts have a sweet taste, for which reason the element was originally called glucinium. It is also the most chemically toxic element. The OSHA permissible working exposure is limited to 2 micrograms/m^3 of air for an 8 hour shift, and about 100 micrograms/month. NIOSH recommendeds that airborne exposure never exceed 0.5 micrograms/m^3. Soluble compounds in the form of solutions, dry dust, or fumes may produce dermatitis or, when inhaled, acute potentially lethal effects similar to those caused by the poison gas phosgene.

It has many outstanding structural properties. Beryllium has a high strength per unit weight. Its modulus of elasticity exceeds that of steel (i.e. it is very stiff). It has the highest melting point of the light metals, and it has excellent thermal conductivity. It is also nonmagnetic and nonsparking. These factors combine to make it a superb material for heat shielding in ICBM reentry vehicles, and for inertial guidance system gyroscopes. It has good resistance to chemical attack (it does not oxidize at ordinary temperatures, and it resists nitric acid). It is very valuable as an alloying agent with copper. Beryllium-copper alloys (2% Be) have high electrical and thermal conductivity, high strength and hardness, good corrosion and fatigue resistance,

and nonmagnetic properties. It is a superior material for springs and is used in many types of machinery, and for electronic connectors. About 75% of U.S. beryllium demand is for use in this alloy.

Its oxide has a very high melting point, excellent thermal conductivity, and high hardness (beryllium's ability to scratch glass is probably due to the formation of a thin layer of the oxide) and strength, making it useful in high temperature ceramics and as a substrate for high-density electronic circuits. It is also transparent to microwaves and used in many microwave applications. About 15% of U.S. beryllium consumption is for the oxide.

Beryl (which contains about 4% Be) is the most important beryllium ore throughout the world. Emerald, one of the rarest and most valuable gemstones, is a form of beryl. The major U.S. source has been low-grade bertrandite deposits in Utah. The United States is one of only three countries that process beryllium ore and concentrates into beryllium products and supplies most of the rest of the world with these products. U.S. beryllium production is exclusively handled by Brush-Wellman Inc's Delta Utah plant. It processes bertrandite and imported beryl. Beryllium hydroxide is shipped to the company's plant in Elmore, OH, where it is converted into beryllium alloys, oxide, and metal. One other company in the United States has the capability to produce beryllium alloys -- NGK Metals Corp., a subsidiary of NGK Insulators of Japan -- which produces beryllium alloys at a plant near Reading, PA. NGK Metals purchases its berylloium from Brush-Wellman however.

Beryllium compounds are generally white (or colorless in solution) and show great similarity in chemical properties to the corresponding compounds of aluminum. This similarity makes it difficult to separate beryllium from the aluminum that is almost always present in beryllium ores. Beryllium metal is usually prepared by reducing the fluoride with magnesium.

In 1997 U.S. beryllium consumption of 205 tons was valued at more than $70 million. Worlwide production in that year was 276 tonnes (beryllium metal content). The use of beryllium (as an alloy, metal, and oxide) in electronic and electrical components, and aerospace and defense applications accounted for more than 80% of consumption. The U.S. price of beryllium as domestic vacuum-cast metal ingot was $720/kg. Beryllium-copper master alloy was $350 per kilogram of contained beryllium, and beryllium oxide powder was $169/kg. The National Defense Stockpile (NDS) goal for beryllium metal was 363 tonnes.

It has many application in nuclear technology. It has high transparency to x-rays. It also has a remarkable range of desirable qualities for neutron-related applications. It generates neutrons when bombarded with gamma rays or alpha particles (about 30 neutrons/million alphas), through the He-4 + Be-9 -> C-12 + n reaction, which made it the first known source of free neutrons for laboratory study (the neutron energy is also a remarkably high 5.7 MeV). The existence of the alpha->n reaction made it useful in the neutron initiator in early fission weapons.

It has a very low absorption cross section for thermal neutrons and is thus used as a canning material for fuel elements in reactors. Due its low atomic weight it is an excellent neutron

moderator, and has the largest macroscopic neutron scattering cross section of any material (due in part to its packing density). It also undergoes Be-9 + n -> Be-8 + 2n reactions, thus acting as a neutron multiplier; although this endothermic reaction (- 2.4 MeV) only becomes significant around 3 MeV.

These characteristics have made it a reflector/moderator of choice for extremely small, light nuclear reactors; and a neutron reflector of choice for low yield, very light nuclear weapons. The production of high quality beryllium forgings suitable for use as a weapon reflector requires the use of hot vacuum isostatic pressing.

6.4.2 Polonium (Po)

Polonium, element 84, was the first element discovered by Mme. Marie Sklodowska Curie in 1898, while seeking the cause of radioactivity of pitchblende from Joachimsthal, Bohemia. It is named after her native country Poland. It was extracted from the uranium ore pitchblende, where it is found in the concentration 0.1 ppm.

Polonium has more known isotopes than any other element with 27 (all radioactive). Po-210 is the third longest lived (the longest lived isotope is Po-209 with a 109 year half-life). The isotope Po-210 (the only readily accessible isotope) is naturally formed from the decay of radium:

Ra-226 -> (1600 y, alpha) -> Rn-222

Rn-222 -> (3.823 day, alpha) -> Po-218

Po-218 -> (3.05 min, alpha) -> Pb-214

Pb-214 -> (26.8 min, beta) -> Bi-214

Bi-214 -> (19.7 min, alpha) -> Tl-210

Tl-210 -> (1.3 min, beta) -> Pb-210 ("Radium D")

Pb-210 -> (21 yr, beta) -> Bi-210 ("Radium-E")

Bi-210 -> (5.01 day, beta) -> Po-210 ("Radium F")

Radium samples more than 30 yr old (to allow the buildup of the Pb-210 parent) contain 0.013% (1 g per 7.5 kg of radium). Polonium can be conveniently extracted from radium by allowing it to deposit on to bismuth plates suspended in RaCl solution. Lead extracted from uranium ores contains Pb-210, from which Po-210 can be periodically harvested. In 1934 it was discovered that it could be manufactured by bombarding natural bismuth (Bi-209) with neutrons to form Bi-210.

Elemental polonium is silvery and has two allotropic forms, the alpha and beta forms, which can coexist between 18 C and 54 C. Its physical properties are:

Density: 9.196 (alpha form); 9.398(beta form); mixtures have intermediate densities

Melting point: 254 degrees C

Boiling point: 962 degrees C

Chemically polonium resembles bismuth and tellurium. Polonium dissolves readily in dilute acids, but is only slightly soluble in alkalis. Due to its very short half-life polonium compounds tend to decompose quickly, particularly organic compounds. Metallic polonium can be prepared from polonium hydroxide.

Despite its fairly high boiling point, polonium is rather volatile - 50% of a sample will vaporize in 45 hrs at 55 C. Part of its volatility can be attributed to its radioactivity. When an atom decays, the recoil from the energetic alpha emission dislodges many atoms. Since 0.50% of it decays each day, a significant fraction of the atoms are ejected on a daily basis.

Po-210 is an alpha emitter with a half-life of 138.39 days. It is notable in that it emits almost no gamma rays, a property that made it important to the Manhattan Project and the early post-war weapon program of the U.S. By encapsulating Po-210 in a specially shaped capsule with beryllium, separated by a thin metal foil or film, a neutron initiator for fission bombs could be made. Turbulence created by sudden compression would cause the Po-210 and beryllium to rapidly mix, and produce neutrons by the Be-9 + alpha -> Be-8 + n + alpha reaction. The foil prevented the poorly penetrating alpha particles from creating neutrons prematurely. Other alpha emitters produce gamma rays in addition, which easily penetrate the foil and knock neutrons loose prematurely.

Po-210 is highly toxic and dangerous to handle. Its short half-life gives it an extremely intense radioactivity of 4490 Ci/g (4500 times as intense as radium-226, 65000 times as intense as Pu-239), its volatility makes it easy to ingest. The maximum allowable body burden is 0.03 microcuries (6.7x10^12 g).

The Manhattan Project found it impossible to prevent people who handled polonium from absorbing it. Fortunately absorbed polonium is not deposited in bone, but is excreted rather rapidly. Monitoring polonium content in urine, and frequent staff rotation, minimized the hazard. On the other hand, polonium deposited on inhaled particles is retained by the lung like plutonium, and can be extremely hazardous in this form. Much of the lung cancer hazard of tobacco smoking actually comes from polonium deposited on smoke particles, not chemical carcinogens.

In the chemical isolation from ore, the ore is treated with hydrochloric acid, and the resulting solution is heated with hydrogen sulfide to precipitate polonium monosulfide, PoS, along with other metal sulfides, such as that of bismuth, Bi2S3, which resembles polonium monosulfide closely in chemical behaviour, though it is less soluble. Because of the difference in solubility, repeated partial precipitation of the mixture of sulfides concentrates the polonium in the more soluble fraction, while the bismuth accumulates in the less soluble portions. The difference in

solubility is small, however, and the process must be repeated many times to achieve a complete separation. Purification is accomplished by electrolytic deposition.

Po-210 produces heat at a rate of 140 w/g. A capsule containing 500 mg will reach 500 C. A milligram or so makes an easily visible blue glow from excitation of the air. Polonium has been used in devices that ionize the air to eliminate accumulation of electrostatic charges. In most applications (such as textile mills) this use has been replaced by safer beta emitters, but it has remained in use in applications where products are very sensitive to radiation exposure especially photographic film plants.

8.2.2 Design and Testing of the First Fusion Weapons

In January 1951, Ulam broke the barrier to progress by inventing the idea of staging: using the energy released by an atomic bomb primary to compress an external fuel capsule. He initially developed the idea as a means to create improved fission bombs, the second stage being a mass of fissionable material. By late in the month he realized that the powerful compression that was possible would overcome the obstacles to efficient large scale fusion reactions. By multiple staging, bombs of virtually unlimited size could be created.

This key idea was not sufficient by itself. Before a workable design could be developed a scheme was needed for generating efficient compression using this energy flux, as was a means for igniting the fuel once it was compressed. Ulam's idea was to use the neutron flux or the hydrodynamic shock wave of the expanding bomb core to achieve compression. Working with Ulam, Teller added additional refinements to this insight during the month of February. Teller's principal contribution during this period was realizing that the thermal radiation flux from the primary was a more promising means of generating the necessary implosive forces. On 9 March 1951, Ulam and Teller jointly wrote a report, _On Heterocatalytic Detonations I. Hydrodynamic Lenses and Radiation Mirrors_, that summarized these ideas.

From this point on Teller increasingly began to claim exclusive credit for the breakThrough, and eventually came to deny that Ulam had made any original or significant contribution.

Later in March Teller added an important additional element to the radiation implosion scheme. Adapting Ulam's idea to use staged implosion to trigger a fission reaction, Teller suggested placing a fissile mass in the center of the fusion fuel. The convergent shock wave would compress this to supercriticality upon arriving at the center, making it act as a "spark plug" to ignite the fusion reaction. This idea is perhaps not strictly necessary, the convergent shock wave will generate very high temperatures in the center any way and might suffice to initiate fusion as it does in modern laboratory inertial confinement fusion experiments.

Since the continuing compression on the fusion fuel would act to confine the fission spark plug, this final combined design concept was termed the "equilibrium thermonuclear". Teller wrote this idea up in a report on 4 April, 1951.

It was only in April 1951 that the necessary physical principles were in hand to allow the development and testing of an actual hydrogen bomb to go forward. More computations were required to design the device than for any other project in human history up to this point (made possible by the recent invention of the programmable computer). The elapsed time from this point until the detonation of the Mike device was less than 19 months, an achievement as remarkable in its own way as the Manhattan Project.

In April 1951 experiments with fusion reactions and atomic bombs were already being prepared by the US as part of the Greenhouse test series, including a test of the idea of fusion boosting. The Greenhouse George test in particular provided a valuable opportunity to evaluate the Teller-Ulam ideas by allowing the observation of radiation effects in heating and compressing (although not imploding) an external mass of fusion fuel.

Since there are several known designs for incorporating fusion reactions into weapons we come to a question that is largely a matter of definition: Which design qualifies as a *true* hydrogen bomb? I will not try to debate this issue here (see Section 11: Questions and Answers), instead I am including descriptions of all of the significant tests that lead to the development and deployment of early thermonuclear weapons.

The tests are listed in chronological order. Each is followed by a brief discussion of its significance to weapons development.

Greenhouse George

Detonated 5/9/51 at 0930 (local time) on a 200 ft tower on Ebireru/Ruby island at Eniwetok atoll.

Total yield: 225 kT

George was a test of a pure fission bomb, and the highest yield bomb tested up to that time. The bomb was a cylindrical implosion U-235 bomb, perhaps based on a design by physicist George Gamow. An experiment called the Cylinder device was piggy-backed on George to test the ignition of a thermonuclear reaction. The cylindrical implosion design allowed the fusion experiment to be heated directly by the pit without the shielding effects of a high explosive layer, and avoided disruption by expanding detonation gases. A deuterium-tritium mixture external to the large fission core was ignited by the thermal flux, and produced detectable fusion neutrons. This was the first ignition of a thermonuclear reaction by an atomic bomb. Rhodes, in _Dark Sun_, estimates the fusion energy yield as 25 kT, although the mass of fuel given ("less than an ounce") is at a factor of 12 too small for this. Other sources simply give the fusion yield as "small". This approach provided no prospect for development into a high yield thermonuclear weapon. This test fortuitously provided useful data for evaluating the Teller-Ulam design which had been devised two months prior.

Greenhouse Item

Detonated 5/25/51 at 0617 (local time) on a 300 ft tower on Engebi/Janet island at Eniwetok atoll.

Total yield: 45.5 kT

First test of a boosted fission device. A deuterium-tritium mixture in the U-235 bomb core boosted fission yield by 100% over its expected unboosted yield. This innovation was eventually incorporated into most or all strategic weapons, but the fusion yield was negligible and overall yield was still limited by the capabilities of fission designs.

Ivy Mike

Detonated 11/1/52 at 0714:59.4 +/- 0.2 sec (local time) at ground level on Elugelab/Flora island at Enewetak atoll.

Total yield: 10.4 megatons.

This was the first test of the Teller-Ulam (or Ulam-Teller) configuration. The Mike device used liquid deuterium as the fusion fuel. It was a massive laboratory apparatus installed on Elugelab Island in the Enewetak Atoll consisting of a cylinder about 20 feet high (more exactly 243.625 inches or 6.19 m), 6 ft 8 in wide, and weighing 164,000 lb (including attached diagnostic instruments); also said to weigh 140,000 lb without "the cryogenic unit" (this may mean the casing by itself). It was housed in an open hanger-like structure 88 ft x 46 ft, and 61 ft high, where assembly started in September of 1952.

The Mike device consisted of a massive steel cylinder with rounded ends, a TX-5 implosion bomb at one end acted as the primary, and a giant stainless steel dewar (thermos) flask holding several hundred liters of liquid deuterium surrounded by a massive natural uranium pusher/tamper constituted the secondary fusion stage (know as the "Sausage").

The welded steel casing was lined with a layer of lead. A layer of polyethylene several centimeters thick was attached to the lead with copper nails. This layer of plastic generated plasma pressure during the implosion.

The Sausage consisted of a triple-walled stainless steel dewar. The inner most wall contained the liquid deuterium. Between this wall and the middle wall was a vacuum to prevent heat conduction. Between the middle wall and the outer wall was another vacuum, and a liquid nitrogen-cooled thermal radiation shield made of copper.

To reduce thermal radiation leakage even further, the uranium pusher (which was oxidized to a purple-black color, making it an excellent thermal radiator) was lined with gold leaf.

Down the axis of the dewar, suspended in the liquid deuterium was a plutonium rod that acted as the "spark plug" to ignite the fusion reaction once the compression shock wave arrived at the center. It did not run the entire length of the dewar, but was supported at each end by axial columns. The spark plug was a boosted fission device, it was hollow and was charged with a few grams of tritium/deuterium gas (which of course liquified once the dewar was charged with liquid deuterium).

The Mike device had a conservative design. The external casing was made of steel and was extraordinarily thick (usually described as "a foot thick", but more likely 10 inches to be consistent with the weight) to maximize the confinement of the radiation induced pressure inside. The interior diameter was thus about 60 inches. A very wide radiation channel was provided around the secondary stage to minimize thermal gradients, and to make success less dependent on sophisticated analysis. Due to the low density of liquid deuterium, and the necessity of thermal insulation, the secondary itself was quite voluminous which, when combined with the wide channel between the secondary and the casing led to the 80 inch diameter. The massive casing accounted for most of Mike's weight (about 85%).

The TX-5 device was an experimental version of the implosion system that was also deployed as the Mk-5 fission bomb. It used a 92 point ignition system, that is, 92 detonators and explosive lenses were used to make the spherical imploding shock wave. This allows the formation of the implosion shock wave with a thinner layer of explosive than earlier designs. The TX-5 was designed to use different fission pits to allow variable yields. The highest reported yield for a TX-5 test was Greenhouse Easy at 47 kT on 20 April 1951, with a 2700 lb device. The smaller mass compared with earlier designs kept the temperature higher and allowed thermal radiation to escape more quickly from the primary, thus enhancing the radiation implosion process. If the Easy configuration was used in Mike, then the secondary fusion/primary yield ratio was 50/1. The deployed Mk-5 had an external diameter of 43.75 inches, the TX-5 would have been substantially smaller since it lacked the Mk-5 bomb casing.

Three fuels were considered for Mike: liquid deuterium, deuterated ammonia (ND3), and lithium deuteride. The reason for choosing liquid deuterium for this test was primarily due to two factors: the physics was simpler to study and analyze, and extensive studies had already been conducted over the previous decade on pure deuterium fuel. The desirability of lithium-6 deuteride as a fuel was known, but sufficient Li-6 could not be produced in time to make the November 1952 target date (in fact construction of the first lithium enrichment plant had just begun at the time of the test).

Liquid deuterium produces energy through four reactions:

1. D + T -> He-4 + n + 17.588 MeV
2. D + D -> He-3 + n + 3.268 MeV
3. D + D -> T + p + 4.03 MeV
4. He-3 + D -> He-4 + p + 18.34 MeV

For Mike to function successfully, densities and temperatures in the secondary sufficient to ignite reactions 2 and 3 were required. This requires densities hundreds of times normal, and temperatures in the tens of millions of degrees K (say, 75 g/cm^3 and 3x10^7 K).

Since the reaction cross section of 1 is some 100 times higher than the combined value of 2 and 3 the tritium is burned as fast as it is produced, contributing most of the energy early in the reaction. Reaction 4, on the other hand, requires temperatures exceeding 200 million K before its cross section becomes large enough to contribute significantly. Whether sufficient

temperatures are reached and quantities of He-3 are produced to make 4 a major contributor depends on the combustion efficiency (percentage of fuel burned).

If only reactions 1-3 contribute significantly, corresponding to the combustion of 25% of the deuterium fuel or less, then the energy output is 57 kT/kg. If reaction 4 contributes to the maximum extent, the output is 82.4 kT/kg. The maximum temperature generated by an efficient burn reaches 350 million K.

The fission fraction for Mike was quite high - 77%. The total fusion yield was thus 2.4 megatons, which corresponds to the efficient thermonuclear combustion of 29.1 kg of deuterium (172 liters), or the inefficient combustion of 41.6 kg (249 liters). The total fission yield was 7.9 megatons, the fission of 465 kg of uranium. All but some 50 kT of this was due to fast fission of the uranium secondary stage tamper by fusion neutrons, a 3.3 fold boost.

The amount of deuterium actually present in Mike was no more than 1000 liters, which is the amount of liquid deuterium handled by Operation Ivy. In fact, it was probably substantially less than this since excess LD2 was undoubtedly brought along in case leakage or other losses occurred.

Prior to test, Mike's yield was estimated at 1-10 megatons, with a most likely yield of 5 Mt, but with a remote possibility of yields in the range of 50-90 Mt. The principal uncertainties here would have been the efficiency of the fusion burn, and the efficiency with which the tamper captured neutrons. Both of these factors are strongly influenced by the success of the compression process. The fusion efficiency involved novel and complex physics which could not be calculated reliably even if the degree of compression were known. The physics for determining the efficiency of neutron capture on the other hand were well understood and could be calculated if the conditions could be predicted.

The upper limit estimate provides some insight into the mass of the uranium fusion tamper. Presumably the 90 Mt figure was calculated by assuming complete fusion and fission of all materials in the secondary. If 1000 liters of deuterium were burned with complete efficiency, the yield would be 13.9 Mt. Fission must account for 76.1 Mt, corresponding to a uranium tamper mass of 4475 kg. Lower amounts of deuterium would lead to higher tamper estimates (a ratio of 0.82 kg of U for each liter of LD2).

The detonation of Mike completely obliterated Elugelab, leaving an underwater crater a 6240 feet wide and 164 ft deep in the atoll where an island had once been. Mike created a fireball 3 miles wide; the "mushroom" cloud rose to 57,000 ft in 90 seconds, and topped out in 5 minutes at 135,000 ft - the top of the stratosphere- with a stem eight miles across. The cloud eventually spread to 1000 miles wide, with a stem 30 miles across. 80 million tons of soil were lifted into the air by the blast.

TX-16/EC-16

The Mike design was actually converted into a deliverable weapon, demonstrating that lithium deuteride is not essential to making a usable weapon. The weaponized design, designated the

TX-16, went into engineering development in June 1952 (5 months before the Ivy Mike test). The design eliminated the cryogenic refrigerator, reduced the weight of the tamper, drastically reduced the dimensions and mass of the casing, used a lighter and less powerful primary, and pared the weight in other areas. The expected yield was reduced to 7 Mt. The device was about 60 inches in diameter, 25 ft long, and weighed 30,000 lb. This weapon design would have been filled with liquid deuterium at a cryogenic filling station before take-off, a reservoir in the weapon held sufficient liquid hydrogen to replace boil-off losses during flight. Components for about five of these bombs were built in late 1953, and had reached deplyment by the time of the Castle tests.

A unit of the TX-16, code named Jughead, was slated for proof test detonation on 22 March 1954 as part of the Castle series, prior to its expected deployment as the EC-16 (Emergency Capability) gravity bomb in May 1954. The excellent results with the solid-fueled Shrimp device in the Castle Bravo test on 1 March(see below) resulted in the cancellation of this test, and then of the entire EC-16 program on 2 April 1954.

Soviet Test: Joe 4/RDS-6s

Detonated: 12 August 1953, on a tower at Semipalatinsk in Kazakhstan

Total yield: 400 kilotons

This was the fifth Soviet trest, and first Soviet test of a weapon with substantial yield enhancement from fusion reactions. This bomb (designated RDS-6s) did not employ the Teller-Ulam configuration, instead it used the "Sloika" design invented by Andrei Sakharov and Vitalii Ginzburg. A sloika is a layered Russian pastry, rather like a napoleon, and has thus been translated as "Layer Cake". The design was first invented in the United States by Edward Teller (who called it "Alarm Clock") but it was not developed into a weapon there.

This design is based on a combination of what Sakharov has called the "First and Second Ideas". The First Idea, developed by Sakharov, calls for using a layer of fusion fuel (deuterium and tritium in his original concept) around a fission primary, with an outermost layer of U-238 acting as a fusion tamper. The U-238 tamper confines the fusion fuel so that the radiation-driven shock wave from the fission core can efficiently compress and heat the fusion fuel to the ignition point, while the low conductivity of the fusion tamper prevents heat loss and at the same time yields addition energy from fast fission by the fusion-generated neutrons. The Second Idea, contributed by Ginzburg used lithium-6 deuteride (with some tritium) as the fusion fuel. Being a solid, this is a convenient material for designing a bomb, and it also produces additional tritium from fission neutrons through the Li-6 + n reaction. This establishes a coupled fission -> fusion -> fission chain reaction in the U-238 tamper, with the fusion fuel acting in effect as a neutron accelerator. Larger bombs can be created by placing additional successive layers of Li-6 D and U around the bomb. The device tested in 1953 probalby had two layers.

A small U-235 fission bomb acted as the trigger (about 40 kT). The total yield was 400 kT, and 15-20% of the energy was released by fusion, and 90% due directly or indirectly to the fusion reaction.

A few weeks before the test it was belatedly realized that despite the sparse population of the area around Semipalatinsk, a serious fallout hazard nonetheless existed for tens of thousands of people. The options were to carry out a mass evacuation or delay the test until an air-dropped system could be arranged, which would take at least six months. Rather than delay the test, a hasty evacuation was conducted. [Note: This implies that the Layer Cake was not available as a usable weapon until after Feb. 1954, a time at which the US had actually deployed the EC-14, a megaton-range lithium deuteride fueled Teller-Ulam design. See the Castle Union test below.]

Castle Bravo

Detonated: 1 March 1954 (0645 local time) on reef 2950 ft off of Nam/Charlie island, Bikini Atoll

Total yield: 15 megatons

The Shrimp device detonated in the Bravo test was the first test of a Teller-Ulam configuration bomb fueled with lithium deuteride. This became the standard design for all subsequent hydrogen bombs (including Soviet designs). Shrimp was a cylinder 179.5 in long, and 53.9 in wide, weighing 23,500 lb. The lithium in Shrimp was enriched to a level of 40% Li-6. The predicted yield of this device was only 6 Mt (range 4-6 Mt), but the production of unexpectedly large amounts of tritium through the fast neutron fission of Li-7 boosted the yield to 250% of the predicted value, making it the largest bomb ever tested by the US (and destroying much of the measuring equipment). The fission yield was 10 Mt, the fusion yield was 5 Mt for a fusion fraction of 33%.

The explosion created a 6000 ft crater, 240 ft deep in the atoll reef. The cloud top rose to 114,000 ft.

The Bravo test created the worst radiological disaster in US history. Due to failure to postpone the test following unfavorable changes in the weather, combined with the unexpectedly high yield, the Marshallese Islanders on Rongerik, Rongelap, Ailinginae, and Utirik atolls were blanketed with the fallout plume. They were evacuated on March 3 but 64 Marshallese received doses of 175 R. In addition, the Japanese fishing vessel Daigo Fukuryu Maru (Fifth Lucky Dragon) was also heavily contaminated, with the 23 crewmen received exposures of 300 R (one later died from complications). The entire Bikini Atoll was contaminated to varying degrees, and many operation Castle personnel were subsequently over-exposed as a result. After this test the exclusion zone around the Castle tests was increased to 570,000 square miles, a circle 850 miles across (for comparison this is equal to about 1% of the entire Earth's land area).

The two stage device Shrimp design was used as the basis for the Mk-21 bomb. The weaponization effort began on 26 March, only three weeks after Bravo. By mid April the military characteristics were defined. On 1 July an expedited schedule for deployment was approved. The use of the final fast fission stage was apparently eliminated. After a number of efforts to

reduce the weight, the design seems to have stabilized in mid-July 1955 with a projected yield of 4 megatons (subsequently tested at 4.5 megatons in Redwing Navajo, 95% fusion, 11 July 1956). Quantity production began in December 1955 and ended in July 1956 with 275 units being produced. The Mk 21 weighed about 15,000 lb; it was 12.5 ft long, and 56 in. in diameter. During June-November 1957 it was converted to the Mk 36 design.

Castle Romeo

Detonated 27 March 1954 on barge in Bikini atoll lagoon near Bravo test site at 0630:00.4 (local time).

Total yield: 11 megatons

The Runt I device (the second in the Castle series) was another solid fueled two stage design. This device was 224.9 in. long, 61.4 in. in diameter, and weighed 39,600 lb. The fuel for Runt was natural lithium deuteride, a major advantage considering the high cost of lithium-6 enrichment. It exceeded its predicted yield by an even larger margin than Bravo, with a most probable yield of 4 Mt out of a 1.5-7 Mt range. This is consistent with the higher proportion of Li-7, compared to Bravo. The fission yield was 7 Mt, for a fusion fraction of 36%.

The Runt I and Runt II devices (seen Castle Yankee below) were design tests for the EC-17 and EC-24 bombs respectively. These two weapons were very similar (externally identical, similar internal configurations, but with different primaries). They were the most powerful weapons ever built by the US, with predicted yields of 15-20 megatons, and were also the largest and heaviest bombs ever deployed by the US The Mk 17/24 (as the deployed versions were eventually designated) was 24 ft. 8 in. long, with a 61.4 in. diameter, and a weight of 41,400-42,000 lb (30,000 lb of this was the 3.5 in. steel casing).

Although the initial work on these weapons dates at least to Feb. 1953, they went into development engineering in Oct. 1953. The EC-17 and EC-24 became the second and third models of hydrogen bomb to enter the US arsenal. From April to September in 1954 EC-17 and EC-24 bombs were stockpiled (5 EC-17, and 10 EC-24). These bombs were removed in October, modified for better safety features and with drogue parachutes for slower fall, and returned to duty as the Mk 17 Mod 0 and Mk 24 Mod 0 in November 1954. These weapons went through two subsequent modifications, and stockpiles reached 200 Mk 17s and 105 Mk 24s during the October 1954 - November 1955 production run. The Mk 24s were retired in Sept-Oct 1956; the Mk 17s were retired between Nov. 1956 and Aug. 1957.

Castle Union

Detonated 26 March 1954 (0610:00.7 local time) on barge in Bikini atoll lagoon off Yurochi Island.

Total yield: 6.9 megatons

This was actually the fourth test in the Bravo series (the third test - Koon - failed when the fusion stage did not ignite). This was the test of the EC-14 Alarm Clock bomb (unrelated to Teller's earlier Alarm Clock concept), which was the first hydrogen bomb actually to enter the US arsenal and the first Teller-Ulam bomb ever to be deployed anywhere. This was a solid fueled two stage device using 95% enriched lithium-6. It also exceeded expectations, the predicted yield was 3-4 Mt (range 1-6 Mt). The fission yield was 5 Mt for a fusion fraction of 28%. The tested device had a length of 151 in., a diameter of 61.4 in,. and weighed 27,700 lb.

The TX-14 Alarm Clock went into development engineering in August 1952, and procurement was approved in mid-September (some 6 weeks before Mike had even been tested). The first EC-14 weapons were produced in Feb. 1954, two months prior to test of the design. The design was simple but had very poor safety features. A total of 5 were deployed, this low figure can probably be attributed to scarcity of Li-6 at the time. Safety could presumably have been improved through retrofitting, but the high cost of these weapons probably led to their rapid retirement. They were removed from the arsenal in October with the deployment of the EC-17. The Mk 14 (its final deployed designation) had a diameter of 61.4 in., a length of 18 ft. 6 in., and weighed 28,954 lbs. After refitting with a drogue parachute its weight increased to 29,851 lbs.

Castle Yankee

Detonated 5 May 1954 (0610:00.1 local time) on barge in Bikini atoll lagoon, above the Union crater.

Total yield: 13.5 megatons

The Runt II device was very similar to Runt I, mostly differing in the design of the primaries. The fuel for Runt II was also natural lithium deuteride. It also exceeded its predicted yield, with a most probable yield of 8 Mt out of a 6-10 Mt range. See Castle Union for a discussion of weapons derived from this test. This device was 225 in. long, 61 in. in diameter, and weighed 39,600 lb. The fission yield was 7 Mt for a fusion fraction of 48%.

Soviet Test No. 19 Test 11/22/55 (No common name) Detonated 11/22/55

Total yield: 1.6 megatons

The first Soviet test of a Teller-Ulam/Sakharov Third Idea bomb. It used radiation implosion to detonate a lithium deuteride fueled capsule. This was the world's first air-dropped fusion bomb test. After this test the Soviet Union used radiation implosion bombs as the basis for their strategic arsenal. Exploded underneath an inversion layer, the refracted shock did unexpected collateral damage, killing three people.