Basics of Thermonuclear fusion

We will now have the first of two lectures about nuclear fusion. In this lecture we will discuss the basic principles of fusion, and in the next lecture we will talk about some specific fusion reactions that are important for stars.

We'll begin by considering a very simplified picture of fusion, which is represented by the cartoon in Figure 1. This figure shows the total energy between two nuclei, as a function of separation in units of fermis (1 fm=10^{-13} cm, which is about the radius of a proton). Remember that there is no *absolute* scale of energy; we need to define a zero of energy, and then energy *differences* are what matter. Here, the zero of energy is when the particles are motionless at infinite separation.

Because both of the nuclei are positively charged, they repel each other electrostatically, and thus the energy increases with decreasing separation r like $E \sim 1/r$. However, when the two get close enough to each other, the attraction of the strong nuclear force means that the energy becomes negative. As a result, if the nuclei get close enough to each other, then they can fuse.

How can this happen? Classically, the only possibility is for the nuclei to have enough energy relative to each other that they can get over the hump in the potential; it would be like rolling a ball over the curve. For a typical reaction, e.g., between two protons, the barrier is about 1 MeV in energy. Thus, classically, the particle has to have that much energy to get close enough to fuse. Expressed in terms of temperature, 1 MeV is about 10^{10} K. Given that for a particular temperature there are some particles more energetic and some less energetic than the average, we could get by with a smaller temperature and still *some* particles would have enough energy to get over the barrier. But even after doing our best, we'd still need a temperature of at least hundreds of millions of Kelvin for ordinary hydrogen to helium fusion. You may know that in reality the Sun's core is about 15 million Kelvin. So there must be something else that is going on.

That something is quantum mechanics, which intriguingly means that quantum mechanics, which is the physics of the very small, has a critical influence on the temperature (and size, based on the virial theorem) of stars! One useful way to picture the role of quantum mechanics is that, because every particle is considered to have a wavefunction with a finite extent, what you really need for fusion is for the wavefunctions of the two nuclei to overlap sufficiently that there is a reasonable probability that they are close enough for fusion. This is thus in contrast to the classical picture that each nucleus is a point.

This leads to one of the many bizarre, but experimentally-verified, consequences of quantum mechanics: a particle can *tunnel* through a potential barrier despite not having the energy to, classically, get over that barrier. For fusion, this means that the energy of the incoming nucleus can be significantly less than 1 MeV, which in turn means that the temperature can be a lot less than a few hundred million Kelvin. See Figure 2 for a representation of the tunneling process.

Mathematically, the transmission probability of a particle of energy E and rest mass m through a potential barrier $V(x)$ (where x is the coordinate of motion) is

$$T(E) = e^{-2\int dx \sqrt{\frac{2m}{\hbar^2}[V(x)-E]}} , \tag{1}$$

where the integral is over the range of x where $V(x) > E$. Any time we see an equation in physics it is good practice to determine whether it makes sense in different ways. Our first check is always of units: any argument of an exponential must be dimensionless, and this one is. We expect that the larger $V - E$ is, and/or the larger the width of the barrier is, the lower the probability must be, and that is also the case. Something that is not as obvious but is an important part of quantum mechanical intuition is the placement of Planck's constant $\hbar$. As $\hbar \to 0$ we see that no matter how thin the barrier is or no matter how little V exceeds E, the transmission probability goes to zero. Indeed, $\hbar \to 0$ is the classical limit.

A summary of this part of the lecture is that fusion can be considered to be a possible consequence of the overlap of the wavefunctions of nuclei. In stars, the way the overlap happens is that the nuclei are moving fast because of their high temperature, and they can therefore push close enough together for fusion. But there is another possibility. If the density is high enough, then (1) the nuclei are already close to each other, and (2) as we will now discuss briefly and will revisit when we talk about white dwarfs and neutron stars, the high density introduces a quantum mechanical energy that applies even if the temperature is zero.

We will approach that energy using the uncertainty principle. We know that, formally, the product of the uncertainty Δx in position with the uncertainty Δp in x-momentum has a lower limit: $\Delta x \Delta p \geq \hbar/2$. But we also know from square well problems that if a particle is confined to a region of size Δx, then the minimum quantum mechanical momentum is in fact $p_{\rm min} \sim \hbar/\Delta x$. This is called the Fermi momentum, p_F, after the Italian physicist Enrico Fermi. If the number density is n then $\Delta x \sim n^{-1/3}$ and thus $p_F \sim \hbar n^{1/3}$. There is an associated energy, the Fermi energy E_F: if we ignore the rest mass-energy mc^2, then the Fermi energy is $E_F = p_F^2/(2m)$ for nonrelativistic momenta $p_F \ll mc$, and $E_F = p_F c$ for ultrarelativistic momenta $p_F \gg mc$. The general formula including the rest mass-energy is the usual one: $E_F = \sqrt{p_F^2 c^2 + m^2 c^4}$. In the same way that high temperature can produce nuclei with high enough energy to fuse, so can high density. The first realm is called *thermonuclear* fusion and the second is called *pycnonuclear* fusion.

Now let's go through some nomenclature regarding fusion reactions.

A reaction of the form

$$a + X \to Y + b \tag{2}$$

is often represented by $X(a, b)Y$. An intermediate state of an excited nucleus is represented with an asterisk: Z^*, or a second type of excited nucleus could be Z^{**}. There might be such an intermediate state, as in

$$p +^{11} B \to^{12} C^* \; . \tag{3}$$

Here the number ^{11} in front of B means the atomic weight (this is roughly the sum of the number of protons and the number of neutrons in the nucleus). The excited state, or compound state, often has many ways in which it can break up. For example, $^{12}C^*$ can break up into $^{12}C^{**} + \gamma$ (a lower excited state plus a photon), $^{11}B + p$ (the original combination of boron and a proton), $^{11}C + n$, and so on. In all cases, when considered in isolation the reactions must conserve (1) energy, (2) linear momentum, (3) charge, (4) baryon number (this is the total number of protons plus neutrons, minus the number of antiprotons plus antineutrons, not that those are common), and (5) lepton number (separately for electron, muon, tau, so the electron lepton number is the number of electrons minus the number of positrons, and so on). Angular momentum must also be conserved, but that's more detailed because we have the intrinsic angular momentum as well as the angular momentum of motion relative to a given center of mass. Remember that when we talk about energy conservation, *all* types of energy have to be included, which means the rest mass energy, the Fermi energy, any binding energy, and so on.

For example, for neutron decay in free space, $n \to p + e^- + \nu$, where the neutrino ν must be an electron antineutrino, to conserve lepton number.

As another example, the reaction $\gamma \to \gamma + \gamma$ in free space satisfies all of these conservation laws, but can it happen? No, it is kinematically forbidden because the two photons must be precisely parallel to conserve momentum, and therefore there is zero solid angle in which this can take place.

For each decay mode i, we can define a mean lifetime for decay τ_i, and for a given lifetime we can estimate the energy width of the decay from the uncertainty principle: $\Gamma_i \approx \hbar/\tau_i$ in energy. We assume that the probability of a decay through this channel over a very short time $\Delta t \ll \tau_i$ is $\Delta t/\tau_i$, *regardless* of how long the nucleus has been in existence. This implies that if there are no other decay channels, then the probability that a decay will occur after a time t is $1 - \exp(-t/\tau_i)$, which reduces to t/τ_i when $t \ll \tau_i$. The assumption that the probability of a decay in a short time Δt is independent of the age or previous history of the nucleus gives us the Poisson distribution, and it is well verified by experiment. This is how we can estimate half-lives of billions of years in some cases: for example, if you have 10^{20} ^{238}U nuclei, which have a half-life of 4.5×10^9 years, then after 1 second you expect $0.5 \times 10^{20} \times (1 \text{ second})/4.5 \times 10^9$ yr ≈ 350 decays. This decay law does not apply to everything, of course; if an average human lifetime is 80 years it does not mean that there is a 1/80 chance of dying every year, otherwise you'd have many people living for centuries!

If there are multiple decay channels, then the probability of decay through a particular channel i is just $(1/\tau_i)/\sum_j(1/\tau_j) = \tau/\tau_i$, where $\tau = 1/\sum_j(1/\tau_j)$ is the total mean lifetime (please note that τ is less than any of the individual τ_j values, as it must be). This can therefore also be written $P_i = \Gamma_i/\Gamma$. This type of *harmonic averaging* (weighting by the reciprocal) will also come up when we discuss opacities.

Cross sections.—In everyday life, the cross section of something is often simply its geometrical area. If you throw projectiles at me, the ones that hit will hit in my area in profile, maybe 0.6 m^2 face-on, and half that edge-on. As long as the projectile is small compared to me, its nature doesn't matter much; a water balloon or a rotten tomato will have the same effective cross section of interaction. If we wanted to, we could define my cross section in another way. Suppose there were to be a steady pelting by rotten tomatoes, but that everyone had such bad aim that any area had the same probability of being hit as any other area. We can then define the tomato flux, $F_{\rm tom}$, as the number per area per second. My cross section is then $R_{\rm tom}/F_{\rm tom}$, where $R_{\rm tom}$ is the rate at which they hit me in number per second. The cross section is not always independent of the nature of the projectiles. For example, a pane of glass is nearly transparent to visible light, but it is opaque to infrared, which is what allows greenhouses to be hot.

For some reaction $X + \alpha \rightarrow Y + \beta$, we therefore define the cross section $\sigma_{\alpha\beta}(v)$, where v is the relative speed between X and α, as the number of reactions per second per target divided by the incident flux of the projectiles. Then the reaction rate per volume is $\sigma_{\alpha\beta}(v)vn_\alpha n_X$ cm^{-3} s^{-1}. If the target and projectile are the same, we need to divide by two to avoid double-counting. Because in general the cross section depends on the relative velocities, to get the average rate we need to integrate $\sigma_{\alpha\beta}v$ over the velocity distribution of the nuclei. In thermal equilibrium, the velocities are in a Maxwell-Boltzmann distribution, and thus so are the relative velocities in the center of mass frame. The velocity-averaged cross section at a temperature T is then:

$$\langle\sigma v\rangle_{\alpha\beta} = \left(\frac{8}{\pi m}\right)^{1/2}(kT)^{-3/2}\int_0^\infty \sigma_{\alpha\beta}(E)e^{-E/kT}E\,dE\ . \tag{4}$$

The cross section can also be written in the form

$$\sigma_{\alpha\beta}(E) = \pi\lambda\!\!\!^{-}{}^2 g\frac{\Gamma_\alpha\Gamma_\beta}{\Gamma^2}f(E)\ . \tag{5}$$

Here $\lambda\!\!\!^{-}$ is the reduced DeBroglie wavelength: $\pi\lambda\!\!\!^{-}{}^2 = \pi\hbar^2/(2Em)$. This, roughly, is what the cross section would be for a perfect absorber. g is a statistical factor. $\Gamma_\alpha\Gamma_\beta/\Gamma^2$ is the probability that the initial particles will get together multiplied by the probability that the end products will be as specified. The final factor, $f(E)$, is the "shape factor", which depends on the process: it can be "resonant" or "non-resonant". Basically, resonant interactions are peaked at some energy, whereas non-resonant interactions are either intrinsically non-resonant or occur at energies well below the resonant energy. We now discuss the non-resonant possibility in more detail, and leave resonances for later.

Non-resonant: For all else equal, we expect the probability of barrier penetration to decrease with increasing electric charge, because then the Coulomb repulsion is greater. Similarly, we expect that for all else equal it should be easier to penetrate a barrier for a higher-energy (or higher-speed) particle. Indeed, the barrier penetration factor is $P_l(E) \propto e^{-2\pi\eta}$, where $\eta = Z_\alpha Z_X e^2/(\hbar v) = 0.1574 Z_\alpha Z_X (\mu/E)^{1/2}$ (μ is the reduced mass, $\mu = m_\alpha m_X/(m_\alpha + m_X)$), in units of the mass of a proton. The cross section for non-resonant interactions can thus be written

$$\sigma_{\alpha\beta}(v) = \frac{S(E)}{E} e^{-2\pi\eta} \ , \tag{6}$$

and the net result is that

$$\langle \sigma v \rangle \propto \int_0^\infty S(E) \exp\left[-\left(\frac{E}{kT} + \frac{b}{E^{1/2}}\right)\right] dE \ , \tag{7}$$

where $b \equiv 0.99 Z_\alpha Z_X \mu^{1/2}$ (MeV)$^{1/2}$. In this expression we see that there is a tradeoff between barrier penetration (increasing with increasing energy) and number of available particles (decreasing with increasing energy). The maximum of the integrand is called the "Gamow peak": $E_0 = 1.22(Z_\alpha^2 Z_X^2 \mu T_6^2)^{1/3}$ keV, after the Russian-American physicist George Gamow. The full width at $1/e$ max is $\Delta \approx 2.3(E_0 kT)^{1/2}$. For example, at 2.2×10^7 K and using the $p+^{12}$C reaction, $E_0 = 30.8$ keV and $\Delta = 17.6$ keV. This reaction is extremely sensitive to temperature: for example, the peak at 2.2×10^7 K is $5\times$ the peak at 2.0×10^7 K!

The shape factor $S(E)$ has to be determined experimentally, but over the energy range of interest it is close to constant. Integration of this expression isn't possible in closed form, however one trick is to integrate a Gaussian that is of the same height and curvature at maximum, and this works reasonably well because the integrand is strongly peaked.

The reactions in the Sun are extremely rare at its core temperature. How do we know? Because the Sun has lasted for billions of years! If the reactions were common then the Sun would already have used up all of its core hydrogen. Given that the reactions are rare, what does this mean about the sensitivity of the fusion reactions to the temperature? Since the reactions have exponential dependences, there is *tremendous* sensitivity to temperature. Thus, a small change in the temperature can affect the energy generation rate a great deal. As a result, the central temperature of a main sequence star, in which ordinary hydrogen is fused into ^{4}He, is very insensitive to the mass of the star.

These reactions are inherently exponential in the temperature, but it is analytically convenient to express things in power law form. It is therefore common to write the energy generation rate per mass ϵ as $\epsilon = \epsilon_0 \rho^\lambda T^\nu$ for density ρ and temperature T.

Now let's think about the nuclei themselves. If we can bring nuclei together, will that result in the net release of energy?

Like many things, the structure of nuclei is determined by what gives a minimization of energy. What are the contributions to that energy? Of the ones that we will consider, the strong force gives a negative (binding) contribution and the electrostatic force gives a positive contribution. But the two have different characters. The most important is that although at very small separations between nuclei the strong force is, well, stronger than the electrostatic force (by two orders of magnitude), the electrostatic force has a $1/r^2$ character whereas the strong force dies off exponentially with separation, with a characteristic range of around a fermi (10^{-13} cm).

We can therefore think about what will happen if we add more nucleons to a nucleus. Starting off with small nuclei, adding nucleons increases the binding energy and thus makes the nucleus more stable. When the nucleus gets large enough, however, the electrostatic repulsion works through the whole volume, whereas the strong force is only local, so electrostatic repulsion starts to win and the binding energy per nucleon decreases (it does not become positive, but it decreases). We see this in the nuclear binding curve shown in Figure 3. This means that we can release energy by fusion if the nuclei are small, or by fission (or other processes such as alpha-particle emission) if the nuclei are large. One other insight that follows from these considerations is that although low-mass nuclei tend to have equal numbers of neutrons and protons (because nuclear binding is then maximized), the most stable high-mass nuclei have more neutrons than protons because fewer protons means less Coulomb repulsion. Nuclei also have shells similar to electron shells. These special numbers of nucleons, which fill the shells, are called "nuclear magic numbers" (2, 8, 20, 28, 50, 82, and 126), and nuclei with that many neutrons and/or protons are extra stable (meaning that they are particularly tightly bound), which is the reason for some of the up-and-down structure that we see in Figure 3.

A closer look at Figure 3 also tells us something very important about nuclear fusion in stars. If you can get from ^{1}H (i.e., a proton) to ^{4}He, then you have released a bit more than 7 MeV per nucleon. Going *all the way to* ^{56}Fe releases just about another 1.7 MeV. That means that the overwhelming majority of the nuclear energy released in fusion is released in the first step, of hydrogen to helium. Moreover, because fusing up the chain (say, helium to carbon or carbon to oxygen) involves nuclei with larger and larger electric charges, greater temperatures are needed and thus the fusion rate is increased. This is why the main sequence phase of a star's life is much longer than the subsequent stages of nuclear burning: at later stages there is less energy available and that energy is used up more quickly.

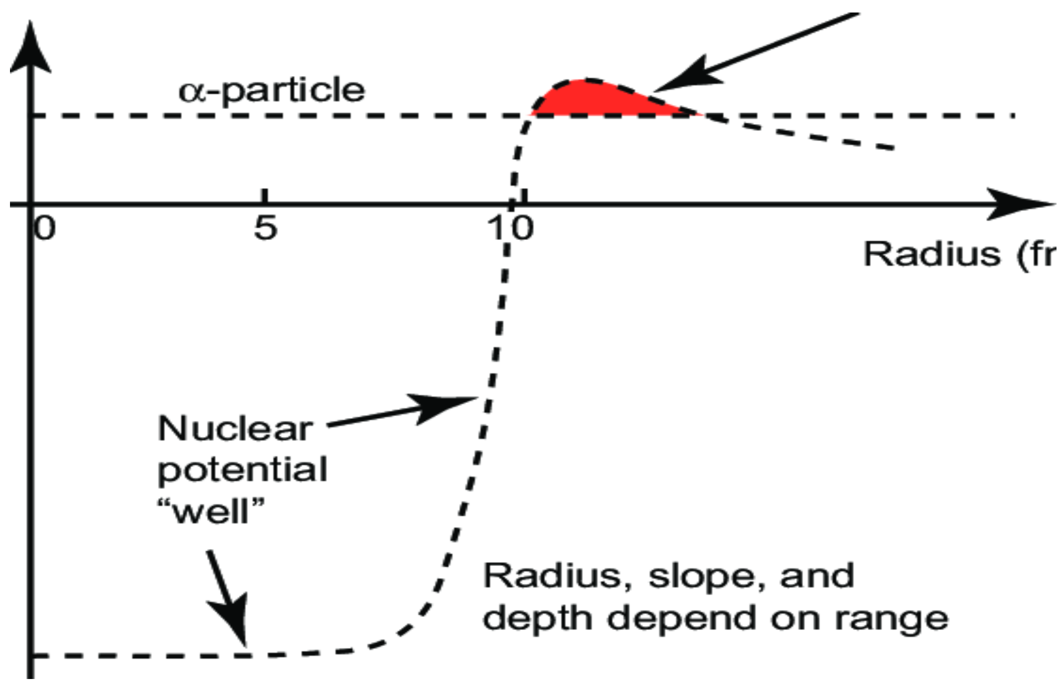

Fig. 1.— Cartoon of the energy potential between a nucleus and an alpha particle (two protons and two neutrons; the nucleus of ^{4}He), from ResearchGate. The radius is in fermis (1 fm=10^{-13} cm, about the radius of a proton). At large distances, electrostatic repulsion means that it takes more and more energy to get closer and closer. But close enough, the attractive strong nuclear force drops the potential to negative energies, which means that energy can be released by the nuclear reaction. However, for this to happen the wavefunction of the alpha particle (in this case) has to overlap enough with the wavefunction of the nucleus for the strong interaction to have a chance. In thermonuclear fusion this happens because the particle has enough energy to at least partially overcome the electrostatic barrier. In pycnonuclear fusion this happens because the density is high enough that the particle starts out close to the nucleus.

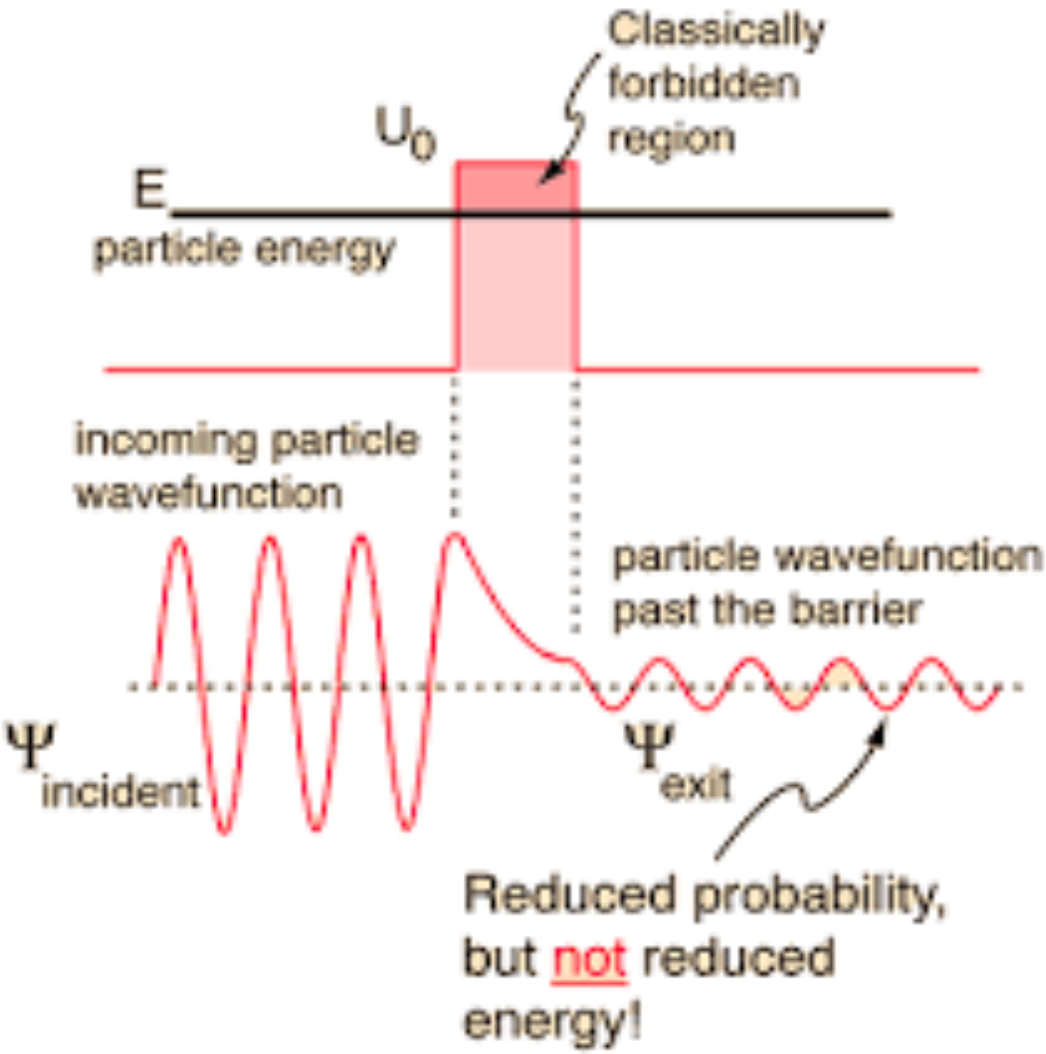

Fig. 2.— Representation of quantum tunneling, from hyperphysics. The particle comes in from the left, with an energy E. It encounters a potential barrier $U_0 > E$; in classical physics the particle would not be able to go through the barrier. However, as we see in the bottom part of the figure, if we consider the particle quantum mechanically then the situation is different. The amplitude (and thus the probability, which is proportional to the square of the amplitude) of the particle's wavefunction suffers an exponential decrease as it goes through the barrier. However, the amplitude does not go to zero. As a result, there is a chance that the particle will get through the barrier. This plays a major role in stellar nuclear fusion because it decreases the required core temperature by more than an order of magnitude.

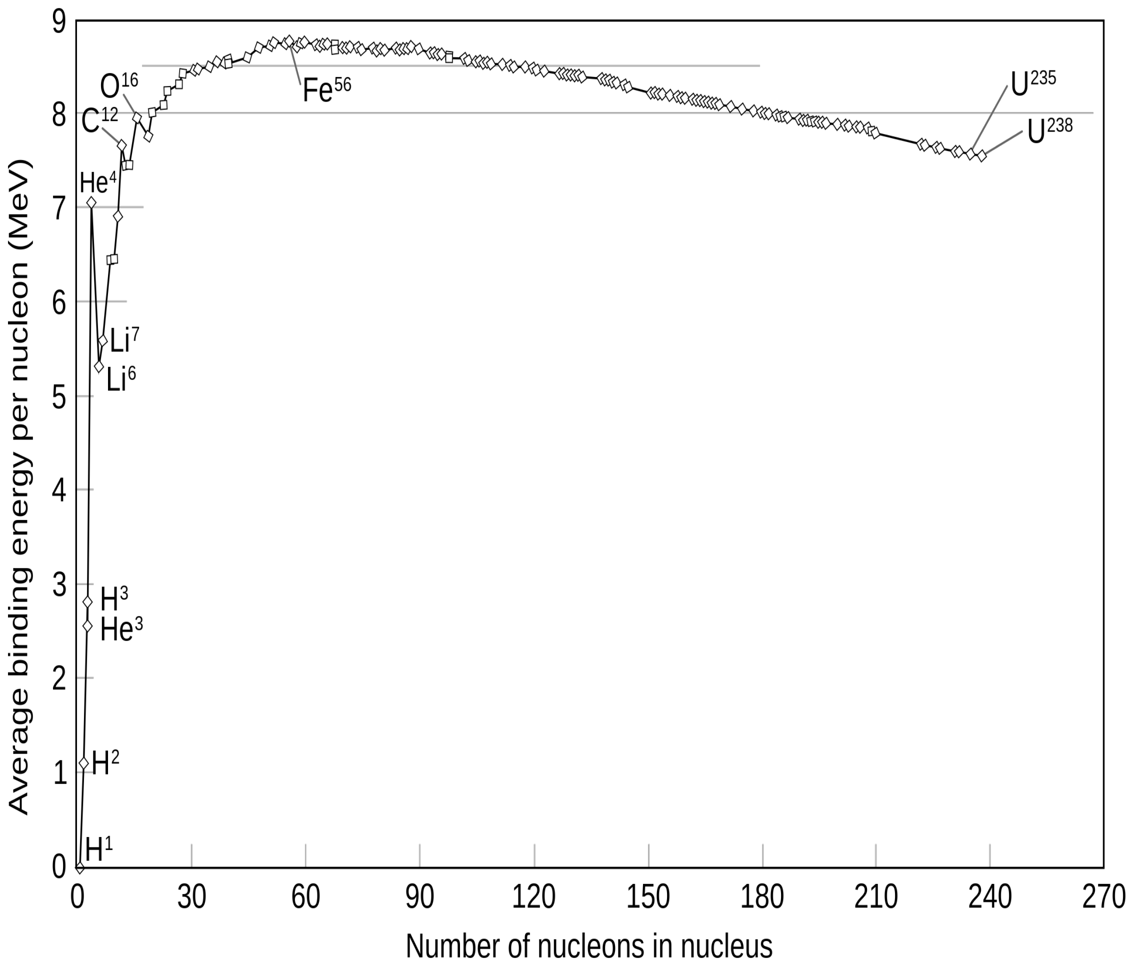

Fig. 3.— Binding energy per nucleon, in MeV, as a function of the number of nucleons in the nucleus (also known as the atomic weight). The normalization is that a single proton (ordinary hydrogen) has zero binding energy. We see that the binding energy increases with atomic weight to iron-group elements, peaking at ^{56}Fe, and then deceases as the atomic weight increases. This means that fusion of light elements can release energy, and fission of heavy elements can release energy.