We left off at the end of the last class with a massive star collapsing on a neutronised core, which continues compressing freely until neutron degeneracy pressure and a stiffening of the equation of state cause it to bounce. We now seek to understand what happens to the remainder of the star. A preview: it blows up.

I. Core collapse supernova physics

We begin our discussion by considering the interior of a star that is about to explode, and asking about the energy budget provided by the central engine, the collapsing iron core.

A. Energy budget

As we have seen, the collapse of the iron core of a massive star occurs on a dynamical timescale. The initial iron core is of order $5,000 - 10,000$ km in radius, and the mass is of order a Chandrasekhar mass, about 1.5 $M_\odot$, so the dynamical time is

$$t_{\rm dyn} \sim \frac{1}{\sqrt{G\rho}} \sim 1 \text{ second.} \tag{1}$$

Thus the core collapses on a timescale that is tiny compared to the dynamical time of the star as a whole – the outer envelope of the star just sits there while the core collapses. The collapse releases a huge amount of gravitational potential energy at the centre of the envelope. Given an initial core mass of $M_c \approx 1.5\ M_\odot$, and an initial radius $R_c \approx 10^4$ km, and a final neutron core of comparable mass and a radius of $R_{nc} \approx 20$ km, the amount of energy released is

$$\Delta E_{\rm grav} \approx -GM_c^2 \left(\frac{1}{R_c} - \frac{1}{R_{nc}}\right) \approx -\frac{GM_c^2}{R_{nc}} \approx 3 \times 10^{53} \text{ erg.} \tag{2}$$

Of this, the amount that is used to convert the protons and electrons to neutrons is a small fraction. Each conversion ultimately uses up about 7 MeV, so the total nuclear energy absorption is

$$\Delta E_{\rm nuc} = 7 \text{ MeV} \frac{M_c}{m_{\rm H}} \approx 2 \times 10^{52} \text{ erg} \approx \frac{\Delta E_{\rm grav}}{15}. \tag{3}$$

Thus only $\sim 10\%$ percent of the energy is used up in converting protons to neutrons. The rest is available to power an explosion.

To explode the star, we must first eject the stellar envelope. The binding energy of the envelope to the core is roughly

$$\Delta E_{\rm bind} = \frac{GM_c(M - M_c)}{R_c} \approx 5 \times 10^{51} \text{ erg} \approx \frac{\Delta E_{\rm grav}}{60}. \tag{4}$$

Thus only a few percent of the available energy is required to unbind the envelope.

The remaining energy is available to give the envelope a large velocity, to produce radiation, and to drive nuclear reactions in the envelope. We don't have a good first-principles theory capable of telling us how this energy is divided up, but we can infer from observations. The observed speed of the ejecta is around 10,000 km s^{-1}, so the energy required to power this is

$$\Delta E_{\rm kin} = \frac{1}{2}(M - M_c)v^2 \approx 10^{51} \text{ erg} \approx \frac{\Delta E_{\rm grav}}{300}. \quad (5)$$

Finally, the observed amount that is released as light is comparable to that released in kinetic energy:

$$\Delta E_{\rm rad} \approx 10^{51} \text{ erg} \approx \frac{\Delta E_{\rm grav}}{300}. \quad (6)$$

Both of these constitute only about 1% of the total power.

So where does the rest of the energy go? The answer is that it is radiated away too, but as neutrinos rather than photons. The neutrinos produces when the protons in the core are converted into neutrons don't escape immediately, but they do eventually escape, and they carry away the great majority of the energy with them. Only about 1% is used to power everything else, but that is more than enough to blow up the star.

Getting that 1% number from first principles, and more generally understanding the mechanism by which the energy released in the core is transferred into the envelope of the star, is one of the major problems in astrophysics today. We have a general outline of what must happen, but really solving the problem is at the forefront of numerical simulation science.

Here's what we know in general outline: as long as the collapsing core has a pressure set by relativistic electrons, its adiabatic index is $\gamma_a = 4/3$. As it approaches nuclear density and more of the electrons convert to neutrons, it initially experiences an attractive nuclear force that pulls it together, and this has the effect of pushing γ_a even lower, toward 1, and accelerating the collapse. Once the densities get even higher, though, the strong nuclear force becomes repulsive, and γ_a increases to a value $\gg 4/3$.

This is sufficient to halt collapse of the core, and from the perspective of the material falling on top of it, it is as if the core suddenly converted from pressureless foam ($\gamma_a < 4/3$) to hard rubber ($\gamma_a > 4/3$). The infall is therefore halted suddenly, and all the kinetic energy of the infalling material is converted to thermal energy. This thermal energy raises the pressure, which then causes the material above the neutron core to re-expand – it "bounces". The bounce launches a shock wave out into the envelope.

The bounce by itself does not appear to be sufficient to explode the star. The shock wave launched by the bounce stalls out before it reaches the stellar surface. However, at the same time all of this is going on, the core is radiating neutrinos like crazy. Every proton that is converted into a neutron leads to emission of a neutrino, and the collapsing star is sufficiently dense that the neutrinos cannot escape. Instead, they deposit their energy inside the star above the core, further heating the material there and raising its pressure. The neutrinos are thought to somehow re-energize the explosion and allow it to finally break out of the star. This is still an unsolved problem. Figure 1 shows a recent example of numerical work trying to solve it.

B. Nucleosythesis

The shock propagating outward through the star from the core heats the gas up to $\sim 5 \times 10^9$ K, and this is hot enough to induce nuclear burning in the envelope. This

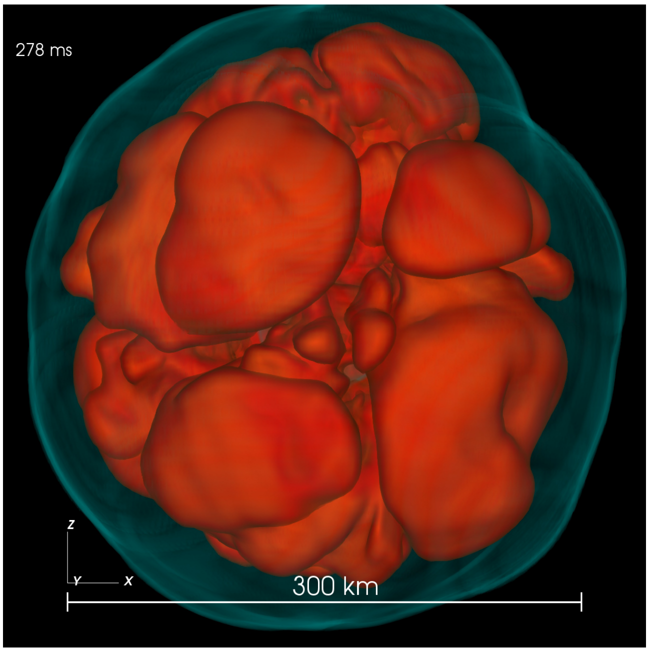

Figure 1: A still image from a 3D simulation of a collapsing stellar core. The volume rendering shows surfaces of constant entropy. The outer, blue surface is the shock where the bouncing material encounters the infalling stellar envelope, and the red mushroom-like features show "boiling" due to neutrino heat deposition. Source: Hanke et al. (2013, *The Astrophysical Journal*, 770, 66, https://arxiv.org/abs/1303.6269. To see the entire simulation, visit https://www.youtube.com/watch?v=2RxIwtxdEnQ.

burning changes the chemical composition of the envelope, creating new elements. Much of the material is heated up enough that it burns to the iron peak, converting yet more of the star into iron-like elements.

I say iron-like because the initial product is not in fact iron. The reason is that the most bound element, ^{56}Fe, consists of 26 protons and 28 neutrons, so it has two more neutrons than protons. The fuel, consisting mostly of elements like ^{4}He, ^{12}C, ^{16}O, ^{28}Si, all have equal numbers of protons and neutrons. Thus there are not enough neutrons around to pair up with all the protons to make ^{56}Fe.

Converting protons to neutrons is via β decays is possible, and in fact it is the first step in the *pp*-chain. However, as we learned studying that reaction, β decays are slow, and in the few seconds that it takes for the shock to propagate through the star, there is not enough time for them to occur.

The net result is that the material burns to as close to the iron peak as it can get given the ratio of protons to neutrons available. This turns out to be ^{56}Ni. This is not a stable nucleus, since it is subject to β decay, but the timescale for decay is much longer than the supernova explosion goes on for, so no beta decays occur until long after the nucleosynthetic process is over.

Not all the material in the star is burned to the iron peak. As the shock wave propagates through the star it slows down and heats things up less. The net results is that material further out in the star gets less burned, so the supernova winds up ejecting a large amounts of other elements as well. Calculating the exact yields from first principles is one of the goals of supernova models.

II. Observable properties

Now that we understand what powers supernovae internally, let us see if we can understand their externally-visible properties.

A. Light Curves

When a supernova goes off, what do we observe from the outside? The first thing,

Figure 2: Observations of supernova shock breakout, made by A. Soderbergh using the SWIFT satellite. The top row shows before the explosion (7 January 2008) and the bottom shows after the explosion (9 January 2008). The left panel is the X-ray image, and the right is the optical image. Source: https://www.nasa.gov/centers/goddard/news/topstory/2008/swift_supernova.html.

which was only seen for the first time in 2008, is a bright ultraviolet flash from the shock breaking out of the stellar surface. We saw this because Alicia Soderberg, then a postdoc at Princeton, got very lucky. She was using an X-ray telescope to study an older supernova in a galaxy, when she saw another one go off. The telescope was observing the star as it exploded, and it saw a flash of X-rays as the shock wave from the deep interior of the star reached the surface. Figure 2 shows the before and after images.

After the initial flash in X-rays, it takes a little while before the optical emission reaches its peak brightness. That is because the expanding material initially has a small area, and most of that emission is at wavelengths shorter than the visible part of the spectrum. As the material expands and cools, its optical luminosity increases, and reaches its peak a few weeks after the explosion. After that it decays. The decay can initially take one of two forms, called linear or plateau, but after a while they all converge to the same slope of luminosity versus time.

This slope can be understood quite simply from nuclear physics. As we mentioned a moment ago, the supernova synthesises large amounts of ^{56}Ni. This nickel is unstable, and it undergoes the β decay reaction

$$^{56}_{28}\mathrm{Ni} \rightarrow {}^{56}_{27}\mathrm{Co} + e^+ + \nu_e + \gamma \tag{7}$$

with a half-life of 6.1 days. This is short enough that most of the nickel decays during the initial period of brightening or shortly thereafter.

However, the resulting ^{56}Co is also unstable, and it too undergoes a β decay reaction:

$$^{56}_{27}\text{Co} \rightarrow {}^{56}_{26}\text{Fe} + e^+ + \nu_e + \gamma. \tag{8}$$

This reaction has a half-life of 77.7 days, and it turns out to be the dominant source of energy for the supernova in the period from a few tens to a few hundreds of days after peak. The expanding material is cooling off, and this would cause the luminosity to drop, but the radioactive decays provide a energy source that keeps the material hot and emitting.

By computing the rate of energy release as a function of time via the β decay of cobalt-56, we can figure out how the luminosity of the supernova should change as a function of time. Radioactive decays are a statistical process, in which during a given interval of time there is a fixed probability that each atom will decay. This implies that the number of cobalt-56 decays per unit time that occur in a particular supernova remnant must be proportional to the number of cobalt-56 atoms present:

$$\frac{dN}{dt} = -\lambda N. \tag{9}$$

Here N is the number of cobalt-56 atoms present and λ is a constant. The equation simply asserts that the rate of change of the number of cobalt-56 atoms at any given time is proportional to the number of atoms present at that time.

This equation is easy to integrate by separation of variables:

$$\frac{dN}{N} = -\lambda\, dt \qquad \Longrightarrow \qquad N = N_0 e^{-\lambda t}, \tag{10}$$

where N_0 is the number of atoms present at time $t = 0$. The quantity λ is known as the decay rate. To see how it is related to the half-life $\tau_{1/2}$, we can just plug in $t = \tau_{1/2}$:

$$\frac{1}{2} N_0 = N_0 e^{-\lambda \tau_{1/2}} \qquad \Longrightarrow \qquad \lambda = \frac{\ln 2}{\tau_{1/2}}. \tag{11}$$

For ^{56}Co, $\lambda = 0.0039$ / day.

While radioactive decay is the dominant energy source, the luminosity is simply proportional to the rate of energy release by radioactive decay, which in turn is proportional to the number of atoms present at any time, i.e. $L \propto N$. This means that the instantaneous luminosity should follow

$$L \propto e^{-\lambda t} \qquad \Longrightarrow \qquad \log L = -(\log e)\lambda t + \text{constant}. \tag{12}$$

Thus for the cobalt-56-powered part of the decay, a plot of $\log L$ versus time should be a straight line with a slope of

$$-(\log e)\lambda = -0.0017\ \text{day}^{-1}. \tag{13}$$

An excellent test for this model was provided by supernova 1987A, which went of in 1987 in the Large Magellanic Cloud, a nearby galaxy. The supernova was observed for more than five years after the explosion, and as a result we got a very good measure of how its luminosity dropped. Figure 3 shows the data. We can see a

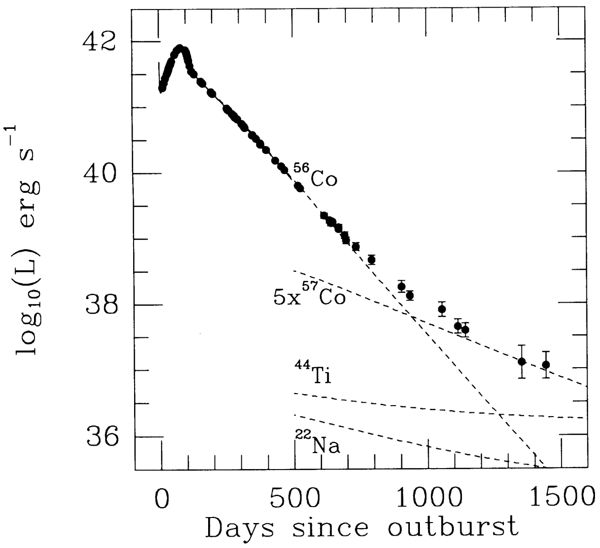

Figure 3: The observed light curve of supernova 1987A, with dashed lines indicating the sources of power at various times. Source: Suntzeff et al. (1992, *The Astrophysical Journal*, 384, L33, http://adsabs.harvard.edu/abs/1992ApJ...384L..33S).

clear period when the slope follows exactly what we have just calculated. Once enough of the ^{56}Co decayed, other radioactive decays with longer half-lives took over. Currently the light output is being powered by decay of ^{44}Ti, which has a half-life of ≈ 60 years.

B. Neutrinos

Supernova 1987A also provided strong evidence for another basic idea in supernova theory: that supernovae involve the neutronisation of large amounts of matter, and with it the production of copious neutrino emission. The first detection of supernova 1987A was *not* its light. The shock wave takes some time to propagate through the star after the core collapses. The neutrinos, however, escape promptly, and on February 23, 1987 three neutrino detectors on Earth detected a burst of neutrinos: the Kamiokande II detector in Japan, the Irvine-Michigan-Brookhaven detector in the US, and the Baksan detector in Russia. The neutrinos arrived more than three hours before the first detection of visible light from the supernova. However, burst is perhaps too strong a word, since the total number of neutrinos detected by all three detectors was 25 – neutrinos are hard to measure!

Nonetheless, this was vastly above the noise level, and provided the first direct evidence that a supernova explosion involves release of neutrinos. It also confirmed the hypothesis that the great majority of the explosion energy must be radiated away in the form of neutrinos, since even from these 25 detections it was possible to roughly compute the neutrino luminosity of the explosion, and compare to the observed optical light output that arrived later.

C. Historical importance

A brief aside: because of their brightness and the long duration for which they are visible, supernovae played an important part in the early development of astronomy, and in the history of science in general. In November of 1572, a supernova went off that was, at its peak, comparable in brightness to the planet Venus. For about two weeks the supernova was visible even during the day. It remained visible to the naked eye until 1574.

The 1572 supernova was so bright that no one could have missed it. One of the people to observe it was the Dane Tycho Brahe, who said "On the 11th day of

November in the evening after sunset, I was contemplating the stars in a clear sky. I noticed that a new and unusual star, surpassing the other stars in brilliancy, was shining almost directly above my head; and since I had, from boyhood, known all the stars of the heavens perfectly, it was quite evident to me that there had never been any star in that place of the sky, even the smallest, to say nothing of a star so conspicuous and bright as this. I was so astonished of this sight that I was not ashamed to doubt the trustworthiness of my own eyes. But when I observed that others, on having the place pointed out to them, could see that there was really a star there, I had no further doubts. A miracle indeed, one that has never been previously seen before our time, in any age since the beginning of the world."

Tycho was so impressed by the event that he wrote a book about it and decided to devote his life to astronomy. He went on to make the observations that were the basis of Kepler's Laws. Kepler himself saw another supernova in 1604. The supernovae played a critical role in the history of science because they provided clear falsification of the idea that the stars were eternal and unchanging, which had dominated Western scientific thought since the time of the ancient Greeks. Previous variable events in the sky, such as comets, were taken to be atmospheric phenomena, and there was no easy way to disprove this. With the supernovae, however, they persisted long enough to make parallax observations possible. The failure to detect a parallax for the supernovae provide without a doubt that they were further away than the Moon, in the supposedly eternal and unchanging realm outside the terrestrial sphere.

III. Thermonuclear supernovae

The supernovae that we have discussed thus far are supernovae driven by the collapse of the core of a massive star. These are ultimately powered by gravity. However, there is another route to supernova explosions, and another energy source: thermonuclear supernovae.

Thermonuclear supernovae occur when a white dwarf that is supported by degeneracy pressure has mass added to it, pushing it above the Chandresekhar mass. There are two scenarios for how this might happen. First, two white dwarfs may find themselves orbiting very closely in a binary system. If they are sufficiently close, they will emit gravitational radiation as they orbit, which will remove energy from the orbit and cause them to eventually collide and merge. This is referred to as the double-degenerate scenario, since it involves two degenerate stars.

The second scenario is that a white dwarf may be in a close binary system with a main sequence star companion. The companion's radius increases as it evolves, and once it is large enough, its outer layers become more bound to the white dwarf than to the core of the parent star. At this point, mass will transfer onto the white dwarf, increasing its mass. This is referred to as the single-degenerate scenario.

Regardless of which astrophysical scenario is responsible for pushing a white dwarf above the Chandrasekhar mass, the result is the same. Once the white dwarf is above the Chandrasekhar mass, it cannot be in hydrostatic equilibrium, and it will begin to collapse. This will raise its temperature until the gas is hot enough to burn the carbon and oxygen to heavier elements.

Something similar happens in the core of a massive star, but in a white dwarf there are two big differences. First, in the core of a massive star, the density is low enough that degeneracy pressure is not dominant, so the burning is under non-degenerate conditions.

On the other hand, the white dwarf is completely degenerate. We learned from our discussion of stability that nuclear burning under degenerate conditions is unstable, because the temperature and pressure are decoupled, so an increase in the burning rate does not lead to a decrease in density or temperature. Instead, it leads to an increase in temperature (and no change in density) that further accelerates burning.

A second difference is that, for a star that still has an envelope, unstable nuclear burning is moderated by the presence of the envelope. Thus for example the helium flash in red giants does not disrupt the star, because it happens in a core that is surrounded by ~ 1 $M_\odot$ of shielding. The core burns until it becomes non-degenerate, and it expands when it does, but it has to do $P\,dV$ work against the heavy envelope. This material absorbs the energy released by the explosion and prevents dramatic changes that are visible from the outside. In a white dwarf, on the other hand, there is no envelope to act as a moderator.

The upshot of all of this is that, when a white dwarf is pushed above the Chandrasekhar mass, the entire star explodes in a huge nuclear detonation. The energy release is comparable to that produced in a core collapse supernova, because in both cases the result is the production of $\sim 1\ M_\odot$ of ^{56}Ni.

In terms of observables, the main difference between this type of supernova and the type we discussed previously is the absence of hydrogen. When a massive star ends its life, the collapsing core is surrounded by an envelope of hydrogen, which produces clear absorption features in the spectrum. When a white dwarf explodes, there is little to no hydrogen around, and hence no hydrogen absorption. Supernovae without hydrogen absorption features in their spectra are called type I supernova, while those with hydrogen are called type II. Massive star supernovae are type II.

It turns out that type I is a bit more complicated, because some core collapse supernovae can be type I as well, if they occur in a star that has lost its hydrogen envelope for some reason, for example because it became a Wolf-Rayet star and blew it off, or because it lost its envelope to a binary companion. The core collapse supernovae are also type I, but they are distinguishable from thermonuclear supernovae in other ways, such as the shape of their light curves. The type I supernovae whose light curves are consistent with being thermonuclear are called type Ia, while the core collapse supernovae in stars that have lost their envelopes are classified as type Ib or type Ic, depending on their exact features.

IV. Supernova remnants

When a supernova goes off, the ejecta first encounter any circumstellar material that was present before the explosion. In many cases, this means that the ejecta run into a pre-existing stellar wind, that was one its way out at hundreds of km s^{-1} only to be overtaken by the supernova ejecta travelling at $\sim 10,000$ km s^{-1}. However, eventually the supernova will encounter the interstellar medium, the diffuse gas that fills the space between the stars. This encounter produces observationally-spectacular structures known as supernova remnants, as shown in Figure 4.

To understand the structure and evolution of supernova remnants, we will consider the idealised problem of an explosion in a uniform medium. The explosion involves the ejection of a mass $M_{\rm ej}$ of ejecta, carrying a total kinetic energy E. The medium in which the explosion occurs has a number density of atoms n – since we are dealing with interstellar gas, it is more convenient to deal with number densities than mass densities, which tend to be ridiculously tiny. For reference, a typical mass of ejecta $M_{\rm ej} \approx 1\ M_\odot$, a typical supernova energy is $E \approx 10^{51}$ erg, and a typical interstellar density is $n \approx 1$ cm^{-3}.

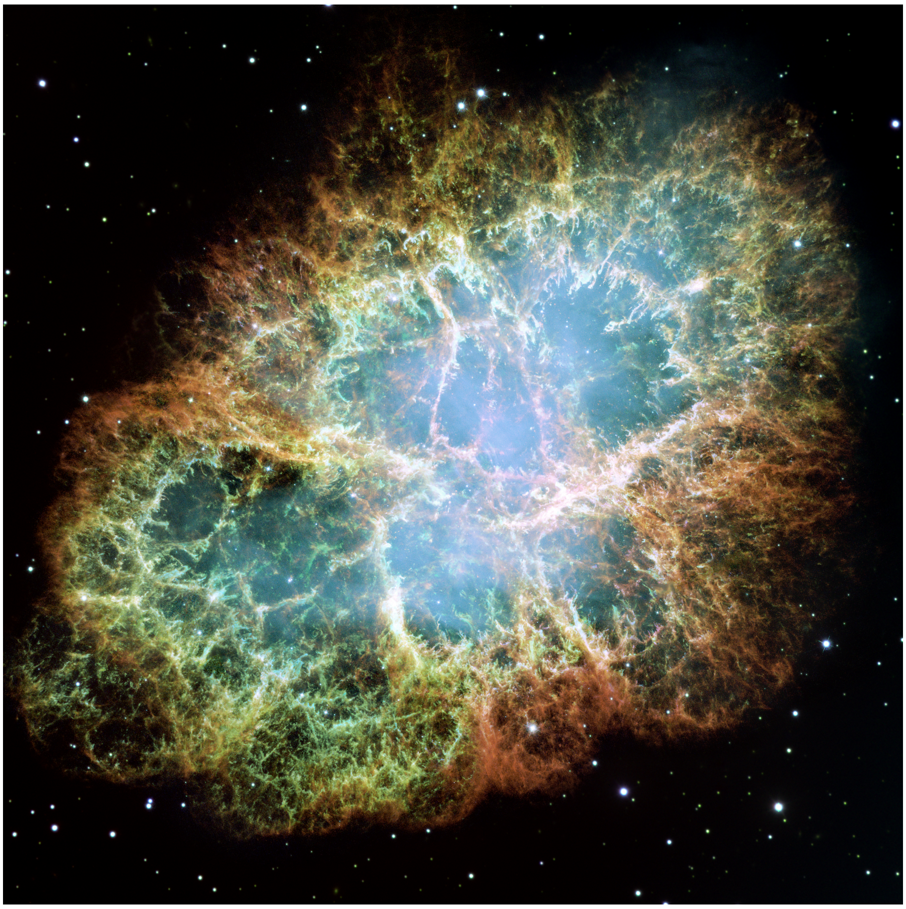

Figure 4: The Crab Nebula, a supernova remnant produced by a supernova that exploded in 1054. We know the date precisely because Chinese astronomers noticed the bright new object in the sky, and recorded it as a “guest star”. Arab astronomers also made records of it. Source: https://en.wikipedia.org/wiki/Crab_Nebula.

A. Free expansion

Immediately after the explosion, the supernova ejecta have only interacted with a mass much smaller than their mass, and so the expansion is essentially free expansion into a vacuum. The expanding ejecta also drive a shock into the surrounding medium, called the forward shock, but it carries negligible momentum initially.

This situation changes when the mass of the swept-up material becomes comparable to the mass of the ejecta, which occurs when the radius of the expanding shell of ejecta is

$$R_1 = \left(\frac{3M_{\mathrm{ej}}}{4\pi n\mu m_{\mathrm{H}}}\right)^{1/3}. \tag{14}$$

Before this point the ejecta travel with mean speed

$$v_{\mathrm{ej}} = \left(\frac{2E}{M_{\mathrm{ej}}}\right)^{1/2}, \tag{15}$$

so the time it takes the ejecta to reach radius R_1 is

$$t_1 \approx \frac{R_1}{v_{\mathrm{ej}}}. \tag{16}$$

Plugging in our numbers from above, the typical radius is $R_1 \approx 2$ pc, the velocity is $v_{\mathrm{ej}} \approx 10^4$ km s^{-1}, and thus the duration of this free expansion phase is $t_1 \approx 200$ yr. Thus supernova remnants do not really being to interact with the interstellar medium for a few centuries after they explode. Before that point the ejecta may glow from the heat of the initial explosion, but will not glow due to running into things.

After the point where the ejecta have swept up about their own mass, the inertia of the ambient material starts to slow down the ejecta. Since the ejecta are traveling highly supersonically, this deceleration launches a shock into the ejecta toward the source star. This is called the reverse shock.

B. The Sedov-Taylor phase

As the reverse shock propagates into the interior of the supernova remnant (SNR), rapidly reaching the centre. At this point all the material in the interior of the SNR is heated to a very high temperature, and its pressure greatly exceeds that of the ambient ISM. We can understand the structure of a supernova remnant at this point using a simple mathematical argument made independently by L. I. Sedov in the USSR and G. I. Taylor in the UK. These authors discovered the solution independently because Taylor discovered it while working in secret on the British atomic bomb project, which was later merged with the American one. It turns out that the problems of a supernova exploding in the interstellar medium and a nuclear bomb exploding in the atmosphere are quite similar physically. Sedov published his solution in 1946, just after the end of World War II, while Taylor's work was still secret.

We are interested in the position of the shock front r as a function of time t. This must depend only on the energy E of the explosion and the density $\rho = n\mu m_{\mathrm{H}}$ of the material being swept up. There are no other parameters to the problem. This simplicity allows us to use dimensional analysis. We have the energy E, density ρ, radius r, and time t, which have units as follows:

$$
\begin{aligned}
[r] &= L && (17)\\
[t] &= T && (18)\\
[\rho] &= ML^{-3} && (19)\\
[E] &= ML^2T^{-2}. && (20)
\end{aligned}
$$

Here L means units of length, T means units of time, and M means units of mass. Thus a density is a mass per unit volume, which is a mass per length cubed. Energy has units of ergs or joules, which is a mass times an acceleration times a distance, and acceleration is distance per time squared.

We want to have a formula for r in terms t, ρ, and E. It is clear, however, that there is only one way to put together t, ρ, and E such that the final answer has the units of length! The mass must cancel out of the problem, so clearly the solution must involve E/ρ. This has units

$$\left[\frac{E}{\rho}\right] = L^5T^{-2}. \tag{21}$$

We want to obtain something with units of length, so clearly the next step is to cancel out the T^{-2} by multiplying by t^2. This gives

$$\left[\frac{E}{\rho}t^2\right] = L^5. \tag{22}$$

Finally, to get something with units of L and not L^5, we must take the 1/5 power. Thus, the radius of the shock as a function of time must, on dimensional grounds, be given by

$$r = Q\left(\frac{E}{\rho}\right)^{1/5} t^{2/5}, \tag{23}$$

where Q is a dimensionless constant. Similarly, the shock velocity as a function of time must follow

$$v = \frac{dr}{dt} = \frac{2}{5}Q\left(\frac{E}{\rho}\right)^{1/5} t^{-3/5}. \tag{24}$$

Actually solving the equations of fluid dynamics shows that

$$Q = \left[\left(\frac{75}{16\pi}\right)\frac{(\gamma_a-1)(\gamma_a+1)^2}{3\gamma_a-1}\right]^{1/5}, \tag{25}$$

where γ_a is the adiabatic index of the gas into which the shock propagates. Taylor used this solution to deduce the energy of the first atomic explosion at Trinity, and you will do the same on your homework.

V. Compact remnants

The final topic of this class is what is left after the supernova. In the case of type Ia supernovae, the answer is "nothing", but core collapse supernovae do leave something behind. We now turn to those things.

A. Neutron stars

First let us ask whether a collapsing iron core will stop collapsing once it neutronises. We said that it will bounce, but can it actually reach hydrostatic equilibrium? We can try to approach this question exactly as we did for white dwarfs, by noting that, once the particles are relativistic, the star will act as an $n = 3$ polytrope and thus have a maximum mass.

Indeed, there is no need to even redo our calculation. We showed that the pressure of a relativistic, degenerate gas with number density n is

$$P = \left(\frac{3}{\pi}\right)^{1/3} \frac{hc}{8} n^{4/3} = \left(\frac{3}{\pi}\right)^{1/3} \frac{hc}{8m_n^{4/3}} \rho^{4/3}. \tag{26}$$

This is exactly the same as the pressure of a degenerate electron gas, except that instead of having $\mu_e m_{\rm H}$ in the denominator of the second expression, we instead have m_n, the mass of a neutron. However, since $m_n \approx m_{\rm H}$, this is exactly the same as for a degenerate electron gas with $\mu_e = 1$. We can therefore use the formula we derived for the Chandrasekhar mass, just with $\mu_e = 1$, and conclude that the maximum mass of a neutron star is

$$M_{\rm max} = \frac{5.83}{\mu_e^2} M_\odot = 5.83\, M_\odot. \tag{27}$$

Unfortunately this turns out to a be a pretty serious overestimate of the maximum neutron star mass, for two reasons. First, this estimate is based on Newtonian physics, and we will show in a moment that the escape velocity from a neutron star is approaching the speed of light. Thus a neutron star has significant general relativistic corrections, which we have not included in our calculation. Second, our calculation of the pressure force neglects the attractive nuclear forces between neutrons; electrons (as long as they are mixed with an equal number of protons) lack any attractive or repulsive forces. The existence of an attractive force reduces the pressure compared to the electron case, which in turn means that only a smaller mass can be supported. How small depends on the attractive force, which is not completely understood. Models that do these two steps correctly suggest a maximum mass of a bit over 2 $M_\odot$, albeit with considerable uncertainty because our understanding of the equation of state of neutronised matter at nuclear densities is far from perfect – this is not an area where we can really do laboratory experiments!

Nonetheless, let us suppose that the core mass is below $\approx 2\ M_\odot$, and thus can be halted in its collapse by neutron degeneracy pressure. What will the properties of the resulting object be? If the neutrons are non-relativistic, then we can treat the star as an $n = 3/2$ polytrope, and use our mass-radius relation for polytropes to find this out. For a non-relativistic degenerate gas of number density n and particle mass m, we have

$$P = \left(\frac{3}{\pi}\right)^{2/3} \frac{h^2}{20m^{5/3}} n^{5/3} = \left(\frac{3}{\pi}\right)^{2/3} \frac{h^2}{20m_n^{8/3}} \rho^{5/3}, \tag{28}$$

where in the second step we used $m = m_n$ and $\rho = m_n n$ for a gas of pure neutrons. Thus the polytropic constant is

$$K_P = \left(\frac{3}{\pi}\right)^{2/3} \frac{h^2}{20m_n^{8/3}}. \tag{29}$$

The mass-radius relation for polytropes is

$$\left[\frac{GM}{-\xi_1^2 (d\Theta/d\xi)_{\xi_1}}\right]^{n-1} \left(\frac{R}{\xi_1}\right)^{3-n} = \frac{\left[(n+1)K_P\right]^n}{4\pi G}, \tag{30}$$

so solving for R we have

$$R = \xi_1 \left\{ \frac{[(n+1)K_P]^n}{4\pi G} \right\}^{1/(3-n} \left[\frac{GM}{-\xi_1^2 (d\Theta/d\xi)_{\xi_1}} \right]^{(1-n)/(3-n)}. \tag{31}$$

Plugging in our expression for K_P, and the values of ξ_1 and $(d\Theta/d\xi)_{\xi_1}$ for $n = 3/2$, we obtain

$$R = 45 \left(\frac{M}{1.4 M_\odot} \right)^{-1/3} \text{ km}. \tag{32}$$

The choice of 1.4 $M_\odot$ is a typical neutron star mass. In reality this radius estimate is a bit on the high side, for exactly the same reasons that our estimate of the maximum mass was too high: we neglected general relativistic corrections, and we neglected attractive forces that reduce the pressure compared to our pure degeneracy estimate. ¡ore sophisticated models give radii closer to ~ 10 km, though again with significant uncertainty coming from the neutron star equation of state.

This radius implies that we do have to consider general relativistic effects: the surface escape speed is

$$v_{\text{esc}} = \sqrt{\frac{2M}{R}} = 0.64c, \tag{33}$$

for $M = 1.4 M_\odot$ and $R = 10$ km.

B. Black holes

If the iron core does exceed the maximum mass that neutron degeneracy pressure can support, there is, as far as we know, nothing that can stop it from collapsing indefinitely. A full description of what happens in such a collapsing star requires general relativity, and is left for the class on that topic. However, we can make some rough estimates of what must happen using general arguments.

As the star collapses, the escape velocity from its surface rises:

$$v_{\text{esc}} = \sqrt{\frac{2GM}{R}}.$$

Once the radius is small enough, this velocity exceeds the speed of light. The critical velocity at which this happens is called the Schwarzschild radius:

$$R_{\text{Sch}} = \frac{2GM}{c^2} \approx 3 \frac{M}{M_\odot} \text{ km}.$$

Thus a neutron star is roughly $2-3$ Schwarzschild radii in size, and it doesn't take much additional compression to push it over the edge.

The Schwarzschild radius is the effective size of the black hole. Nothing that approaches within that distance of the mass can escape, since nothing can move faster than light. Because nothing that happens inside the Schwarzschild radius can ever influence events outside it, the Schwarzschild radius is called an event horizon.

C. Accretion power

How do we observe neutron stars and black holes? The answer is that, when they're all alone, for the most part we don't. A bare black hole is, by definition, completely free of any kind of emission. A bare neutron star does radiate, but only very weakly. It luminosity is

$$L = 4\pi R^2 \sigma T^4.$$

Neutron stars are born very hot, $T > 10^{10}$ K, but after ~ 1 Myr the star cools and the temperature drops to $\sim 10^6$ K. Plugging in $R = 10$ km with that temperature gives $L = 0.2\ L_\odot$. This is dim enough to make it quite hard to detect any but the nearest neutron stars by thermal emission, particularly since, at this temperature, the emission peaks in the x-ray, and must therefore be studied from space. We have indeed identified some of the nearest and youngest neutron stars, such as the one at the centre of the Crab Nebula, by their thermal x-ray emission. However, this is not an option for most neutron stars.

Instead, we tend to see these objects only when they have a companion that donates mass to them. The energetics work as follows. Consider a neutron star of radius R that accretes an amount of mass dM in a time dt. The material falls from rest at infinity, so it has zero energy initially. Just before it arrives at the surface, its potential and kinetic energies must add up to zero, so

$$\frac{1}{2}v^2\,dM - \frac{GM\,dM}{R} = 0 \qquad \Longrightarrow \qquad \frac{1}{2}v^2 = \frac{GM}{R}.$$

When the material hits the surface and stops, its kinetic energy is converted into heat, and then it is radiated away. In steady state all the extra energy must be radiated, so the amount of energy released is

$$dE = \frac{1}{2}v^2\,dM = \frac{GM\,dM}{R}.$$

The resulting luminosity is just the energy per unit time emitted via this process:

$$L_{\mathrm{acc}} = \frac{dE}{dt} = \frac{GM}{R}\dot{M}.$$

In the case of a black hole there is no surface to hit, but infalling material will still usually radiate away energy, because if it has angular momentum it will have to go into orbit in the form of a disc before accreting. As mass moves inward through the disc, it tends to radiate a significant fraction of its gravitational binding energy.

Accretion luminosity increases as the radius of the star decreases, which means that it can be a very potent energy source for compact things like neutron stars and black holes. For example, suppose a star accretes at a rate of $10^{-10}\ M_\odot\ \mathrm{yr}^{-1}$, so that it gains roughly $1\ M_\odot$ of mass over the age of the universe. For the Sun, $L_{\mathrm{acc}} \approx 10^{-3}$ $L_\odot$, unnoticeably small. For a white dwarf, $R = 0.01\ R_\odot$, it would be $L \approx 0.1\ L_\odot$, high enough to be brighter than just an isolated white dwarf normally is. For a neutron star, $R = 10$ km, it would be $100\ L_\odot$, and for a black hole, $R \approx 3$ km, it approaches $1000\ L_\odot$!

Of course this process cannot produce arbitrarily high luminosities, for the same reason that stars cannot have arbitrarily high luminosities: the Eddington limit. The the luminosity is too high, then radiation forces are stronger than gravity material will be pushed away from the accreting object rather than attracted to it. The Eddington limit is

$$L_{\mathrm{Edd}} = \frac{4\pi cGM}{\kappa},$$

and if we require that $L_{\mathrm{acc}} < L_{\mathrm{Edd}}$, then we have

$$\frac{GM}{R}\dot{M} < \frac{4\pi cGM}{\kappa} \qquad \Longrightarrow \qquad \dot{M} < \frac{4\pi cR}{\kappa}.$$

Thus there is a maximum accretion rate onto compact objects. The value of κ that is relevant is usually κ_{es}, since usually the accreting material is hot and fairly low density.