

Supernovae

6

Supernovae

Stellar Catastrophes

1. Observations

Which stars explode? Which collapse? Which outwit the villain gravity and settle down to a quiet old age as a white dwarf? Astrophysicists are beginning to block out answers to these questions. We know that a quiet death eludes some stars. Astronomers observe some stars exploding as supernovae, a sudden brightening by which a single star becomes as bright as an entire galaxy. Estimates of the energy involved in such a process reveal that a major portion of the star, if not the entire star, must be blown to smithereens.

Historical records, particularly the careful data recorded by the Chinese, show that seven or eight supernovae have exploded over the last 2000 years in our portion of the Galaxy. The supernova of 1006 was the brightest ever recorded. One could read by this supernova at night. Astronomers throughout the Middle and Far East observed this event.

The supernova of 1054 is by far the most famous, although this event is clearly not the only so-called "Chinese guest star." This explosion produced the rapidly expanding shell of gas that modern astronomers identify as the Crab nebula. The supernova of 1054 was apparently recorded first by the Japanese and was also clearly mentioned by the Koreans, although the Chinese have the most careful records. There is a strong suspicion that Native Americans recorded the event in rock paintings and perhaps on pottery. An entertaining mystery surrounds the question of why there is no mention of the event in European history. One line of thought is that the church had such a grip on people in the Middle Ages that no one having seen the supernova would have dared voice a difference with the dogma of the immutability of the heavens. One historian, the wife of one of my colleagues, has an interesting alternative viewpoint. She argues that the people who made careful records of goings-on in medieval Europe were the monks in scattered monasteries. Some of these monks were renowned for their drunken revelries and orgies in total disregard for their official vows of abstinence and celibacy. Would such people have shied from making mention of a bright light in the sky when they kept otherwise excellent records? (Never put it in writing?) The truth may be more mundane, having to do with weather or mountains blocking the view. A report of a few years ago called attention to a reputed light in the sky at the time of the appointment of Pope Leo, but this has not been widely accepted. In any case, there is no confirmed record of the supernova of 1054 in European history.

Five hundred years later, the Europeans made up for lost time. The supernova of 1572 was observed by the most famous astronomer of the time, the Danish nobleman Tycho Brahe. Tycho made the careful measurements of planetary motions that allowed his student, Johannes Kepler, to deduce his famous laws of planetary motion. Tycho also carefully recorded the supernova of 1572. His data on the rate at which the supernova brightened and then dimmed in comparison to other stars gives a strong indication of the kind of explosion that occurred. The

heavens favored Kepler in his turn with the explosion of a supernova in 1604. Kepler also took careful data by which we deduce that he witnessed the same kind of explosion as his master. Although there are counterarguments and some controversy, both Tycho's and Kepler's supernovae are widely regarded to be the kind of event modern astronomers label Type Ia.

Shortly after Kepler came Galileo and his telescope and then Newton with his new understanding of the laws of mechanics and gravity. This epoch represented the birth of modern astronomy. Astronomers now have large telescopes, the ability to observe in wavelengths from the radio to gamma rays, and the keen desire to study a supernova close up. Ironically, however, Kepler's was the last supernova to be observed in our Galaxy. Supernovae go off rarely and at random, so a long interval with none is not particularly surprising, just disappointing. We do observe a young expanding gaseous remnant of an exploded star, a powerful emitter of radio radiation known as Cassiopeia A. From the present size and rate of expansion of the remnant, we deduce that the explosion that gave rise to Cas A occurred in about 1667. By rights, this should have been Newton's supernova, but no bright optical outburst was seen. Evidently, this explosion was underluminous. There are reports that Cas A was seen faintly by John Flamsteed, who was appointed the first astronomer royal of England by King Charles II in 1675, but there are questions concerning the timing and whether or not that sighting was in the same position as the remnant observed today. Astrophysicists have calculated that supernovae are brighter if they explode within large red-giant envelopes (see Section 6). The suspicion is that the star that exploded in about 1667 may have ejected a major portion of its envelope before exploding or that the star was otherwise relatively small and compact. That condition, in turn, may have prevented Cas A from reaching the peak brightness characteristic of most supernovae. We will see in Chapter 7 that supernova 1987A, the best-studied supernova of all time, had this property of being intrinsically dimmer than usual.

A new chapter in this story was written by the *Chandra X-Ray Observatory* launched on July 23, 1999. After astronomers had searched for decades with other instruments, the *Chandra Observatory* found the compact object that was demanded to exist in the remnant of this massive star. Ironically, the very first image obtained by *Chandra* for publicity purposes was of Cas A, since everyone knew an image of Cas A would be spectacular. To everyone's surprise and delight, there was a small dot of X-ray emission right in the dead center of the expanding cloud of supernova ejecta. This central source is putting out X-rays with a luminosity of only about one-tenth that of the total light of our Sun. This explains why it was not seen before. The compact object in Cas A is much fainter than the neutron star in the Crab nebula and, as of this writing, we have still not figured out if it is a neutron star operating under its own power or a black hole with a disk feebly emitting while accreting matter from its surroundings. The *Chandra* web site asked for readers to vote between these choices. That is an amusing exercise, perhaps, but it is not the way science is done. One give-away might be a regular pulse of emission, a frequent clue to a rotating neutron star (Chapter 8), but not expected for emission from a disk around a black hole. So far, no such pulsed emission has been seen.

There is a theme that runs through this discussion of historical supernovae and their currently-observed remnants, both compact and extended, but that is not immediately obvious. That is that there are two kinds of explosions, ones that leaves behind compact remnants and ones that do not. Among the latter are SN 1006 and Tycho's supernova of 1572. Among the former are the Crab nebula and Cas A. As we will explore in this chapter, there are two

fundamental explosion mechanisms, one associated with the collapse of the core of a massive star that must leave behind some sort of compact remnant, either a neutron star or black hole, and the other that blows the star, believed to be a white dwarf, to smithereens, leaving no compact object. A more subtle, but significant, clue is that when a compact object is observed, there is evidence for some sort of elongated appearance or even directed, jet-like flow. This is true for both the Crab and Cas A and for other historical supernovae, events recorded by the Chinese in 386 and 1181, and a nearby, well-studied remnant in the constellation Vela that is thought to have exploded about 10,000 years ago. This correlation also applies to SN 1987A (Chapter 7). The opposite seems also to be true; that when no compact remnant is seen, there is no substantial elongation. SN 1006 and Tycho are examples of this. An exception is the remnant of Kepler's supernova of 1604, about which arguments still rage. *Chandra* X-ray images of Kepler show some elongation and so I will bet it is of the core-collapse variety, although there is, as yet, no sign of a compact remnant. We will explore the significance of the association of compact objects and jet-like structure in Section 5.

All supernovae directly observed since 1604 (with the possible exception of Cas A), and hence all supernovae seen by modern astronomers, have been in other galaxies. Any single galaxy hosts a supernova only rarely. Supernovae occur roughly once per 100 years for spiral galaxies like the Milky Way. Astronomers do, however, observe a huge number of galaxies at great distances. The chance that some of these galaxies will have supernovae go off in them is appreciable. Before supernova 1987A, about thirty supernovae were recorded every year. Closer attention was paid to discovering supernovae after supernova 1987A, and the current rate of discovery is about 100 per year. Many of these supernovae are so distant and so faint that scant useful data are obtained from them, but special programs have yielded good data on very distant supernovae. This will be discussed in Chapter 12.

From the studies of supernovae in other galaxies, astronomers have come to recognize that there are two basic types called, cleverly enough, Type I and Type II. This differentiation was first made in the 1930s when Fritz Zwicky began systematic searches for supernovae at Caltech. The categories of supernovae are traditionally defined by the *spectrum* that reveals the composition of the ejected matter. Complementary information is obtained from the *light curve*, the pattern of rapid brightening and slower dimming followed by each event. As more supernovae have been discovered, the dividing lines of this taxonomy have been blurred by events that share some properties of Type I and some of Type II. As for any developing science, one begins with categories and then seeks to replace mere categories with a solid base of physical understanding.

The spectra of *Type I supernovae* are peculiar in that they reveal no detectable hydrogen, the most common element in the Universe. Some Type I supernovae, called Type Ia, appear in all kinds of galaxies – elliptical, spiral, and irregular. Type Ia tend to avoid the arms of spiral galaxies. Because the spiral arms are the site of new star formation, the suggestion is that Type Ia supernovae explode in older, longer-lived stars. This implies that the progenitor stars of Type Ia supernovae are not particularly massive because massive stars live only a short time. Just how low the mass of these Type Ia supernovae may be is a question of current debate. The light curve for Type Ia supernovae is very identifiable. There is an initial rise to a peak that takes about 2 weeks and then a long slower period of gradual decay over time scales of months that is very similar for all these events. The data recorded by Tycho and Kepler suggest that they both

witnessed Type Ia supernovae. No other galactic supernova has sufficient records to make an identification by type. For decades, all Type Ia supernovae were thought to be virtually identical, but more recent careful observations have revealed small, but real, variations among them.

Near the peak of their light output, *Type II supernovae* show normal abundances in their ejected material, including a normal complement of hydrogen. The material observed at this phase is very similar to the outer layers of the Sun. These supernovae have never appeared in elliptical galaxies. Type II supernovae occur occasionally in irregular galaxies, but mostly in spiral galaxies and then within the confines of the spiral arms. The reasonable interpretation is that the stars that make Type II supernovae are born within the spiral arms and live an insufficient time to wander from the site of their birth. Because they are short-lived, the stars that make Type II supernovae must be massive. The light curve of a typical Type II supernova shows a rise to peak brightness in a week or two and then a period of a month or two when the light output is nearly constant. After this time, the luminosity will drop suddenly and then less rapidly with a time scale of months. This pattern of light emission with time is consistent with an explosion in the core of a star with a massive, extended red-giant envelope, as will be explained in Section 6.

To confuse the issue, one and maybe two other varieties of hydrogen deficient supernovae were identified in the 1980s. These are called, with a further flight of imagination, *Type Ib* and *Type Ic*. The two types are probably closely related. Unlike Type Ia, but like Type II, Types Ib and Ic only seem to explode in the arms of spiral galaxies. Therefore, Types Ib and Ic are also associated with massive stars. Type Ib show evidence for helium in the spectrum near maximum light. Type Ic show little or no such evidence for helium. On the other hand, both types show evidence for oxygen, magnesium, and calcium at later times. This is the strongest argument that Types Ib and Ic are closely related. They show little or no evidence for the strong line of silicon that is a major characteristic of the spectra of Type Ia. Type Ia supernovae show essentially only iron at later times, another factor emphasizing their difference from Types Ib and Ic. The composition revealed by Types Ib and Ic is similar to that expected in the core of a massive star that has been stripped of its hydrogen. In the case of Type Ic, most of the helium is gone as well. This suggests an origin in a star much like a Wolf-Rayet star, but a direct connection to this class of stars has not yet been established. The light curves of Types Ib and Ic are somewhat similar to those of Type Ia, but are dimmer at maximum light.

A bright supernova observed in 1993, SN 1993J, gave yet more clues to the diversity of processes that lead to exploding stars. SN 1993J revealed hydrogen in its spectrum, so this event was a variety of Type II. As the explosion proceeded, however, the strength of the hydrogen features diminished, and strong evidence for helium emerged. In this phase, SN 1993J looked much like a Type Ib. There were a few events like this known before, and several have been seen since. Apparently this star had most, but not all, of its hydrogen envelope removed, probably in a binary mass transfer process. In other cases, the removal of hydrogen is more nearly complete, and in yet others, for the Type Ic, the helium is removed, too. There is yet no direct observational proof for binary companions in Types Ib or Ic or the transition events like SN 1993J, but computer models suggest this is the case for SN 1993J, at least. Strong winds from massive stars could play a role for the Types Ib and Ic, and the relative importance of winds versus binary mass transfer has not been resolved.

2. The Fate of Massive Stars

The evidence indicates that Types II and Ib/c supernovae represent the explosion of massive stars. These stars have presumably evolved from the main sequence to red giants and have had a series of nuclear-burning stages producing ever heavier elements in the core. Just which massive stars participate in this process is still debated.

One way to deduce the masses of the stars that make supernovae is to examine the rate at which the events occur in various galaxies. The death rate can then be compared to the rate at which stars are born with various masses. We know that there are many low-mass stars born every year in a galaxy like ours and rather few massive stars (*why* this should be true is a question under active investigation). If we consider stars with mass in excess of about 20 solar masses, we find such stars are born, and hence die, too infrequently to account for the rate at which Type II supernovae explode. If we consider stars with less than about 8 solar masses, we find that such stars die in excess profusion. Stars with mass between about 8 and about 20 solar masses are born and die at the rate of about once per 100 years in our Galaxy. This is also the approximate rate at which we deduce Type II supernovae occur. Type II supernovae probably come from stars of this mass range. Many of these stars, particularly on the upper end of this mass range, are thought to form iron cores that collapse to form neutron stars. There is thus a strong suspicion that Type II supernovae leave neutron stars as compact remnants of the explosion, and that the gravitational energy liberated in forming the neutron star is the driving force of the explosion.

The rate of explosion of Types Ib and Ic supernovae is not well known because relatively few of them have been discovered. Their rate is roughly the same as that for Type II. This suggests that Types Ib and Ic come from roughly the same mass range as Type II. One possibility is that Types Ib and Ic come only from Wolf-Rayet stars that formed by the action of strong stellar winds in stars more massive than 30 solar masses (Chapter 2, Section 2). This is probably not the only source of Type Ib and Ic events. Because very massive stars are rare, there would probably be too few of them to explain the rate of explosions. This suggests that Types Ib and Ic also come from stars that were born with less than 30 solar masses. A binary companion would then be necessary to help strip away the hydrogen envelope. Nevertheless, the basic arguments that pertain to Type II supernovae hold also for Types Ib and Ic. If Types Ib and Ic come from massive stars, to account for their rate of occurrence and their sites in spiral arms, Types Ib and Ic are also very likely to be associated with core collapse to form neutron stars.

At the lower end of the mass range suspected to contribute to Type II supernovae, the evolution may be slightly different. The outcome, core collapse, is basically the same. Computer calculations show that for stars with original mass between about 8 and 12 solar masses the core will be supported by the thermal pressure when carbon is burned. This stage of carbon burning is then regulated and gentle in the standard way. The carbon burns to produce neon and magnesium, but the oxygen that typically coexists with the carbon after helium burning does not get hot enough to burn. As the core, now composed of oxygen, neon, and magnesium, contracts, the quantum pressure comes into play before any other fuel can ignite. The stars in the mass range 8–12 solar masses will therefore form cores supported by the quantum pressure and consisting of oxygen, neon, and magnesium. The atomic nuclei of neon and magnesium are capable of absorbing an electron, thus turning one proton into a neutron, and transmuting themselves into an element of lower proton number. This process reduces the electrons that are responsible for the quantum pressure that is supporting the core. The result is that the core

collapses before any of the elements in the core begin thermonuclear burning. During the collapse, the remaining nuclear fuels – oxygen, neon, and magnesium – are converted to iron. The net result is a collapsing iron core, just as for the more massive stars where the iron core forms before the collapse ensues. These two processes of iron core collapse may give identical results, or there may be some subtle difference between collapse triggered by absorbing electrons rather than by heating and disintegrating the iron. These differences could affect the explosive outcome. There is some evidence that stars in the lower-mass range with the collapsing oxygen/neon/magnesium cores may be especially efficient in producing some of the rare heavy elements like platinum.

A different way of addressing the question of which stars explode is to ask which stars do not explode because they cast off their envelopes gently and leave white-dwarf remnants. This question has been addressed by counting the number of white dwarfs in stellar clusters of various ages and then estimating what stars must have produced those white dwarfs. Such estimates are roughly consistent with the statement that all stars below about 8 solar masses make white dwarfs, and hence do not make supernovae, at least not right away.

Estimates of the rate of formation of neutron stars in the Galaxy are similar to estimates of the rate of formation of Type II supernovae. This does not prove that Type II supernovae produce neutron stars, but the notion that the two processes are directly related is a nearly universal working hypothesis. The problem with this hypothesis is that no calculations have been able to satisfactorily show that the energy liberated in forming a neutron star can routinely cause an explosion. Despite rather gross changes in the physics over the last three decades, many calculations keep stubbornly predicting no explosion, but total collapse. This does not necessarily mean that such explosions do not occur in nature. The calculations may leave out some important piece of physics. That physics might be presently unknown to us, or the process might be too complex to calculate effectively, like the effects of rotation or magnetic fields. Alternatively, we may find that not all stars that develop collapsing cores do form explosions. Some may leave black holes with no explosion at all.

3. Element Factories

Stars with an initial mass larger than 20 solar masses should form iron cores that collapse. There are so few of these stars that whether they explode or not will not change the total supernova rate appreciably. Some other way must be devised to determine whether or not they explode. The observation that suggests that some of the massive stars must explode is the simple but profound one that says that about 1 percent of the material in stars is composed of elements heavier than helium. These elements cannot be produced in the big bang. Alternatively, we know from theoretical calculations that heavy elements in reasonable proportions are produced naturally in the massive stars in the process of forming an iron core. The conclusion is that at least some of the most massive stars must explode in order to eject their complement of heavy elements into space to be incorporated in new stars.

Calculations show that stars with mass between 8 and about 15 solar masses contain too little in heavy elements outside the collapsing core to contribute substantially to the production of elements like carbon, oxygen, and calcium that are abundant in stars as well as in our bodies. Thus the stars that are presumed to account for most, if not all, of the Type II supernovae are not significant contributors to synthesis of the heavy elements. Stars with mass between about 15

and 100 solar masses produce substantial amounts of heavy elements. If these stars explode and eject their heavy elements, this freshly synthesized material will mix with the hydrogen in the interstellar gas. New stars form from this enriched mixture. If all stars from 15 to 100 solar masses explode, the new stars will have about the right amount of all the abundant heavy elements.

This picture has led to the widespread belief that the most massive stars must explode and produce the heavy elements. There is probably a great deal of truth in this notion. As observations get more accurate, however, there are hints that the broad picture must be reassessed. Detailed stellar spectra of both young and old stars have allowed new accurate measurements to be made of the way that various elements have been produced throughout the history of the Galaxy. There is a suggestion that if all the massive stars from 15 to 100 solar masses explode, many of the basic heavy elements like carbon, oxygen, and iron will be produced in greater quantity than is observed in the stars in the vicinity of the Sun. A possible resolution of this dilemma is that some of the massive stars collapse completely. In this picture, some massive stars would explode, ejecting heavy elements and leaving neutron stars behind as compact remnants. Others would produce no explosion and would leave behind black holes as the only remnant of their previous stellar existence.

The pattern that seems to best satisfy all our present knowledge would have stars from about 8 to about 30 solar masses exploding and those from 30 to 100 solar masses collapsing and swallowing all their heavy elements. The most reasonable position is probably to conclude that we do not yet know enough about the nuclear and evolutionary processes in stars to conclude with any certainty which stars explode and eject the heavy elements we see.

4. Collapse and Explosion

In the collapse of an iron core, the protons capture electrons and convert to neutrons. Each reaction creates a neutrino. This is the process by which the composition is converted to neutrons, the necessary step to make a neutron star. For every neutron formed, there must also be a neutrino. The result is a *lot* of neutrinos.

When the collapse reaches the density of atomic nuclei, the strong nuclear force has a repulsive component. This provides a strong outward pressure. In addition, the quantum pressure of the neutrons plays a role. The result of the increased pressure is that the collapse halts (temporarily, at least). The basic processes as they are thought to occur in a massive star are shown in Figure 6.1.

If you drop something heavy, like a bowling ball, appreciable energy is released when it lands. The more massive the object, the greater the energy released. The farther the object falls, the greater will be the energy released. Imagine dropping the bowling ball from the top of a tall building. Imagine dropping a sports utility vehicle from the top of a tall building. Now imagine the gigantic release of energy when a star with the mass of the Sun collapses to the tiny size of a neutron star, only a few kilometers across. A huge energy is released when the neutron star forms. This energy is several hundred times more than is necessary to blow off the outer layers, those containing calcium, oxygen, carbon, and helium, and any outer envelope of never-burned hydrogen. The problem is that most of the energy produced in the collapse is lost to the neutrinos that can easily stream out of the newly born neutron star and through the infalling matter. If 99

percent of the energy is lost, 1 percent can remain. That is enough to cause an explosion. If 99.9 percent is lost, however, that is too much. The explosion will fail, and the outer matter will continue to rain in and crush the neutron star into a black hole.

The exact treatment of this problem has proven to be very difficult. The requirement is to determine whether 99 or 99.9 percent of the energy is lost to neutrinos, or whether it is some fraction in between. The energy lost to neutrinos must be determined to about one part in a thousand. Uncertainties in the complex physics involved in core collapse have been larger than this critical difference. A related problem is that the explosion process tends to be self-limiting. If more of the energy is trapped, then the rate of infall of new matter from the outer parts of the star is slowed. This, however, decreases the rate at which the collapse produces energy that can power the explosion. The result has been that for decades computer calculations have tended to give results that teeter on the edge of success, some giving explosions, many giving complete collapse to form black holes with no explosion. No completely clear, accepted, reproducible, result has emerged. The stars know how to produce these explosions, but astrophysicists are struggling to figure it out.

Over the last couple of decades, research on this topic has involved two basic mechanisms by which the collapse of an iron core might be partially reversed to make a supernova explosion. One is called *core bounce*. When the neutron star first forms, the new star overshoots its equilibrium configuration giving a large compression to the neutron core. There is then a rebound. This rebound sends a strong supersonic shock wave back out through the infalling matter. The core takes about 1 second to collapse after instability sets in. The core bounce creates the shock in about 0.01 second. If everything works, in this short time a huge explosion should be generated.

If the shock wave is sufficiently strong, the outer matter is ejected, and the neutron star is left behind; however, the shock must run uphill into the infalling matter. Some of the energy of the shock is dissipated by the production and loss of neutrinos. The shock also must do the work of breaking down the infalling iron into lighter elements, protons and neutrons, to form the neutron star. The shock wave can thus stall with insufficient energy to reach the outer layers of the star. Matter can continue to rain down on the stalled shock front, as illustrated in the bottom part of Figure 6.2. The shock front hangs in mid-flow, much as a bow wave stands off a rock in the middle of a stream, as shown in the top of Figure 6.2. The matter will continue to be shocked as material hits this front, but the shocked matter will settle onto the neutron star, just as the water will be slowed, but not stopped, by the rock in the stream. When enough matter lands on the neutron star, the neutron star will be crushed into a black hole. Most calculations currently show that the core bounce alone is not sufficient to cause an explosion.

The other mechanism that has been actively considered takes advantage of the tremendous stream of neutrinos leaving the neutron star. Normal matter, the Sun, is essentially transparent to neutrinos because neutrinos interact only through the weak nuclear force. The only exception to this is neutron star matter. This matter, nearly as dense as an atomic nucleus, is so dense that it can be opaque or at least semitransparent to the neutrinos. Although most of the neutrinos will get out into space, a small fraction will be trapped in the hot matter that lies just behind the shock front created by the core bounce. The slow accumulation of neutrino heat may provide the pressure to reinvigorate the shock, driving the shock outward and causing the explosion. Slow in this case means about a second.

The mechanism for depositing a small fraction of the neutrino energy behind the shock may be related to the boiling of the newly formed neutron star, as shown in Figure 6.3. When the collapse is first halted and the neutron star rebounds, the neutron star is very hot. This heat can cause the neutron star to boil much like a pan of water boils on the stove. The boiling provides a mechanism for carrying the heat upward, in the case of the pan, or outward, in the case of the neutron star, by mechanical motion that bodily carries the heat. Under the right circumstances, this boiling process can be much more efficient in transporting heat than a slower process of leaking radiation, or neutrinos. Calculations of this process in neutron stars are very challenging because the motion is complex. All modern calculations that can follow motion in more than one (radial) dimension show that neutron stars do boil. There is a consensus that explosions will not occur without this boiling. There is still debate about whether this process is sufficient to cause an explosion.

5. Polarization and Jets: New Observations and New Concepts

For the past thirty years, most calculations of core collapse and subsequent events treated the configuration as spherically symmetric. Even if the neutron star boils, the structure of the neutron star may, on average, be spherically symmetric. There are a number of lines of evidence, however, that the explosions that result from the core collapse process are intrinsically non-spherical. Matter may be ejected more intensely in some directions than others.

Some hints of this perspective have been with us for a long time. As we will see in Chapter 8, we observe hundreds of neutron stars as rotating, magnetic *pulsars*. If we look at the supernova remnants, the expanding clouds of gas that have produced neutron stars in supernova explosions, there is evidence for non-spherical behavior. The famous Crab nebula is hardly round. X-ray images obtained by the *Chandra Observatory* show a torus of matter shed by the neutron star and jets of high velocity matter being spurted out in opposite directions along the axis of the torus. The neutron star is even running away in space directly along this jet direction. Cas A shows evidence of a jet-like flow in one direction and a somewhat more diffuse, but distinct flow in the opposite direction. Unlike the case for the Crab pulsar and a couple of other examples, the compact object in Cas A seems to be running away perpendicular to the orientation of the jets, not along them. That must be a clue, but we do not yet know what it is telling us about the explosion process and compact object in Cas A.

The situation until recently was thus that we knew that the left-overs of core collapse were frequently rotating, magnetic neutron stars. What we did not know was whether the rotation and magnetic fields were crucial to the process, or present but incidental to the explosion. Similar arguments applied to the supernova remnants that showed evidence for asymmetries of various kinds. Were these aspects of a few peculiar supernovae, or was something systematic going on?

The technique of measuring the *polarization* of the light from supernovae provided a new window of observations and major new insights into the explosion process. Electromagnetic radiation consists of an electric component oscillating in a fashion that is perpendicular to the magnetic component, with both perpendicular to the direction of motion of the electromagnetic wave (or photon of light in the quantum description). The process of measuring the polarization of the light is one of determining the direction in which the electric field is oscillating. In supernovae, the light scatters off electrons in the outer material of the supernovae before

proceeding to astronomer's telescopes, millions of light years away. This scattering gives an average net orientation of the electric component that is perpendicular to the surface of the supernova. If the supernova is perfectly spherical (or if it is at least round in the aspect it presents to us) all directions will be represented in the light and there will be no net direction to the electric component. If, however, the supernova matter is asymmetric in some fashion, then some parts of the matter will provide more light, and more heavily represent the orientation of the surface they represent, than others. The net effect will be to impart a net orientation to the electric component of all the light from the supernova and this will give a net polarization for astronomers to measure. The basic nature of this effect is illustrated in Figure 6.4. The bottom line is that if a supernova reveals a net polarization it *cannot* be spherically symmetric; it might be pancake shaped, or cigar shaped, or, much more likely, some more complicated shape, but it cannot be round.

Starting about a decade ago our group at Texas, led by my colleague Lifan Wang who was then a Hubble Postdoctoral Fellow here, began to collect polarization data on every supernova that was accessible to us. The early days were hard. Lifan used a small telescope and had to add up data from several nights to get enough signal. This was reminiscent of the heroic days of astronomy early in the 20th century when a single night's data was simply not enough. We learned that lesson and have migrated our program, now led by Dietrich Baade at the European Southern Observatory, to the magnificent Very Large Telescope (VLT) array in Chile where similar observations on their eight-meter telescopes can be done in a half hour!

The first thing Lifan noticed as the data began to come in was that there was a distinct difference between Type Ia supernovae and all the core collapse supernovae. Near and after peak light, Type Ia were barely polarized, if at all. They were essentially round. All of the core collapse supernovae showed significant polarization. They were definitely not round! As even more data came in over the last few years we realized that the strength of the polarization got larger as the supernova aged, thinned out, and allowed us to see deeper into its depths. This meant that the cause of the asymmetry was not some incidental aspect of its environment, but that the inner depths were asymmetric; the very machine driving the explosion was severely out of round. We also realized that Type Ic were highly polarized. These supernovae have lost their hydrogen and helium envelopes allowing us to see deeper into the explosion even at early times. The lesson is the same. The inner depths, driven by the explosion process, are highly non-spherical.

Another important lesson was that in many of the cases, the polarization was not random. The net orientation of the electric field always pointed in the same direction independent of time or even of the wavelength observed. This meant that the supernova ejecta were somehow driven along a special direction during the explosion. Even more data has shown that this behavior is not universal. Sometimes more than one direction is indicated by different ejected elements and sometimes the data seem to indicate truly random directions in space. Still, this tendency for the ejected supernovae matter to "point" in a special direction is a powerful aspect in many cases and a strong clue to what is going on.

If, in common circumstances, the supernova is somehow "pointing" to a certain direction in space, how can that happen? What would tell an exploding star that one direction was somehow special? The obvious answer seems to be rotation. A sphere at rest will have no special orientation, but a rotating sphere, or a planet like the Earth, or a star like the Sun, or the Galaxy

(which is not spherical) for that matter, have a special direction, the direction aligned with the axis of rotation. Rotation automatically selects a special direction. What that specific direction is depends on the accident of birth and maybe subsequent jostling, but that direction is an intrinsic characteristic of a rotating object. There are, however, ways of setting up special directions that do not require rotation. One is that the new-born neutron star may end up oscillating with respect to the outer stellar material: neutron star to the left, star to the right; then vice versa. We have to keep such possibilities in mind as we go forward.

As the polarization data first began to accumulate, the first thing we thought of were jets. Jets blowing along special directions are a ubiquitous aspect of gravitating accreting systems. Proto-stars blow jets. We see jets of matter coming from the centers of galaxies and from black holes in binary stellar systems (see Chapter 10). The infall of the iron core to form a neutron star is an extreme case of a gravitating, accreting system. Perhaps, we thought, a similar thing was happening in the core collapse supernovae.

Another important ingredient in this context is magnetic fields. As outlined above (and will be explored in detail in Chapter 8), pulsars are neutron stars that both rotate and are magnetic. Most of the theories of how to produce jets depend on tangling up magnetic fields. Perhaps, then, magnetic fields are also important to the actual process of the explosion of the magnetic field. This is hard to prove, but I think my student and colleague Shizuka Akiyama and I have taken an important step in this direction. We have examined the physics of the magneto-rotational instability that was first discussed in Chapter 4 in the context of accretion disks. Amplifying magnetic field by this mechanism requires a gravitating system with *shear*, the process by which some matter slides past other matter. The flow in accretion disks intrinsically involves shear; the matter closer to the central star naturally moves faster than the matter further out. The same thing is true in core collapse. As the iron core collapses to form a neutron star, it is like a skater pulling in his arms (Figure 1.2); the neutron star will spin much faster than the original iron core. The difference naturally forms a shear and drives the magneto-rotational instability that will rapidly grow any feeble magnetic field that might be present in the original iron core. The implication is that the magnetic field will naturally grow in this environment. It is not consistent to consider only rotation and ignore the magnetic field. Rotation and magnetic fields will come hand-in-hand in the core collapse environment. The important issue is just how big is the magnetic field and just what it will do to the matter. This is a tough problem, but, to my mind, the polarization is telling us that rotation and magnetic fields are intrinsically coupled to the explosion process, shaping the explosion if not actually causing it.

The polarization then points to an important role for rotation and magnetic fields in the very explosion process itself. If this is the case, then the current numerical calculations may be missing a major ingredient necessary to yield an explosion. The most obvious mechanism for breaking the spherical symmetry by singling out a specific direction is rotation because rotation defines a rotation axis. Proper treatment of rotation, abetted by magnetic fields, may be necessary to understand fully when and how collapse leads to explosions. All the energy of collapse is provided by gravity. This energy temporarily goes into two components: the hot bath of neutrinos that will slowly leak out of the neutron star and the tremendous fly-wheel of the rotating neutron star itself. Tapping the energy of that fly-wheel and sending it up the rotation axis may be just the process that explodes and shapes core collapse supernovae. Adding the effects of rotation and magnetic fields is even more of a computational challenge, but computer

power grows steadily, and progress will be made in this area in the next few years. Other suggestions that rotation and magnetic fields are important to the core collapse process are presented in Chapter 11.

To pursue the question of the role of jets in supernovae, my colleague Alexei Khokhlov, then at the Naval Research Laboratory now at the University of Chicago, explored what jets might do to supernovae. This calculation glossed over a number of complications that need to be investigated more deeply, but addressed fundamental issues by assuming that a newly-formed neutron star could launch jets along the rotation axes in about a second while the outer parts of the star hovered, waiting to be blasted into space or to collapse into a black hole depending on the outcome of the collapse. To correspond to a Type Ib or Ic supernova, the hydrogen envelope of a massive star model was omitted, and only the core of helium and heavier elements was retained (Khokhlov and my Texas colleague Peter Höflich have since done calculation covering more general conditions). The jets penetrated to the surface of the helium core in about 6 seconds. As they propagated, the jets drove bow shocks that blow sideways as well as forward, much as a motor boat creates a bow wave as it powers across a lake. Unlike a lake, a star is basically spherical and the bow waves blown away from the jet open up away from the jet like a flower petal and wrap around the star. If the jets are basically symmetrical in the “up” and “down” direction, the down-going bow shock from the “up” jet will collide with the up-going bow shock from the “down” jet at the equator. The result is that shortly after the jets penetrate the surface, the sideways bow shocks converge and eject the matter out along the equator. If the star has no hydrogen envelope, as assumed by Khokhlov, then the final result is two jets of matter along the axes and a strongly asymmetric, doughnut-like explosion otherwise, as illustrated schematically in Figure 6.4. This is the generic shape predicted for a jet-induced supernova. Although the polarization observations cannot uniquely prove this is the shape, the data are consistent with this shape as the source of the observed polarization in many cases.

Although I have used various props to illustrate this generic shape of jet-induced supernovae (once a carrot and chocolate doughnut were the only supplies available; very messy), my favorite is a breadstick and a bagel because it alliterates. The breadstick threaded through the hole in the bagel represents the matter ejected in the jets. The bagel represents the matter blown out along the equator by the converging bow shocks. This concoction captures some of the sense of the nature of the explosion, but one must recall that it is a static image; in reality matter will be rushing outward in both the “bagel” and the “breadstick” directions.

The explosion computed by Khokhlov was driven entirely by the jets. The stalled shock and the neutrinos described in Section 4 played no role. This trial calculation does not prove that jets alone explode supernovae, but it does show that sufficiently strong jets can do so in principle. Further study will probably show that both jets and neutrinos are necessary in varying degrees. If jets are a critical part of the explosion in many, if not all, core collapse events, then many issues such as nucleosynthesis and the production of black holes must be reconsidered.

These developments leave open the issue of how jets are formed in supernovae if, indeed, they are. One aspect of the problem is that the magnetic fields probably do not represent the strongest force during the core collapse process; the magnetic forces are intrinsically weaker than the pressure forces in the neutron star. On the other hand, the pressure and gravity basically push along a radial direction and cancel one another. The magnetic field has the special property that it can push laterally where ordinary pressure and gravity offer little resistance. Magnetic fields

may help to direct matter and energy toward the rotation axes. By catalyzing the motion of energy in that direction, magnetic fields may help to tap the rotational energy to flow into axial jets without contributing to the brute energy of the flow itself.

One aspect of this problem that Shizuka Akiyama and I have recently emphasized is the somewhat counter-intuitive notion that the final spin and magnetic field of the neutron star will be an irregular function of the original spin of the iron core. If the iron core spins slowly, the neutron star will spin slowly and generate only a weak magnetic field. If the iron core spins a bit faster, the neutron star will spin a bit faster and generate a stronger magnetic field. If, however, the iron core spins faster than a certain amount, then the centrifugal force of rotation will tend to give an extra source of support to the neutron star in addition to its normal pressure. That means that the neutron star will not collapse quite as far or achieve quite so high a density. That, in turn, means that the neutron star will rotate a bit more *slowly*, like a skater who has only pulled her arms in part way. It also follows that the magnetic field generated will be less strong for this faster iron core rotation.

The rotation of the iron core may thus be an important determinant of the final outcome of the collapse. It is conceivable, for instance, that very slowly rotating iron cores will fail to trigger an explosion (as all the most sophisticated computer calculations today show!). Somewhat faster rotation of the iron core will generate more rotation of the neutron star and stronger magnetic fields, perhaps triggering successful jet-induced (and neutrino-boosted) supernovae. With even higher rotation of the iron core, however, the neutron star rotates less fast, generates weaker magnetic fields and perhaps total collapse to form a black hole ensues in this situation. This is only a hypothesis, but it illustrates how thinking about the core collapse problem might change once rotation and magnetic fields are brought into the picture.

What would make one star have a slower rotating iron core and another a faster rotating core? This is also a difficult problem that is the subject of current active research. The evolution of stars from the main sequence to the iron core phase will tend to be accompanied by a migration of angular momentum outward from the faster inner core to the slower outer envelope, thus slowing the spin of the iron core that ultimately forms. The rate at which the core is spun down may also be a sensitive function of the magnetic field that exists in the star, another focus of current research. In addition, the outcome is probably influenced by whether the star has a binary companion. Two stars in orbit can induce a mutual *torque* on one another, thus pumping some of their orbital energy and angular momentum into the spin of the cores of the stars, yielding, other things being equal, faster spinning cores. In other circumstances, the stars could form a common envelope (Chapter 3, Section 9). The two stars might eject the common envelope and form a new compact binary system, but it might be even more likely that the two stars (or an immersed star and the core of the star that formed the common envelope) merge to form one exceedingly rapidly rotating core that could, if the circumstances are right, proceed to form an iron core. The issue of the success of a supernova and whether a given star yields a neutron star or a black hole might then depend on whether or not the star was born in a multiple star system.

A subject that is developing as this second edition goes to print is our recent recognition that rotating neutron stars will be subject to forming shapes that are not only not spherically symmetric, but not even axially symmetric, shapes like spiral arms and other, more complex geometries. Most of the work showing this behavior has ignored both the fact that a new-born

neutron star will still be immersed in the supernova environment with matter raining down on it and that the neutron star will be magnetic. In this case, we again hypothesize, the non-axially symmetric motion will rattle the magnetic field, generating magnetohydrodynamic waves that will sap the energy of the rotation and carry that energy somewhere else, maybe up the axes in jets. It will take some effort to explore these ideas thoroughly, but we again see the expanding range of possibilities once rotation and magnetic fields are considered.

6. Type Ia Supernovae: The Peculiar Breed

The principal peculiarity of Type I supernovae is that such events have no hydrogen in their ejected material. The hydrogen envelope that surrounds most stars has either been ejected or consumed to helium or heavier elements. As noted in Section 1, there are two rather different observed categories of Type I. Some of them, the Types Ib and Ic, like Type II, occur only in spiral or irregular galaxies. The Type Ia supernovae occur in all types of galaxies. This makes Type Ia events different in some fundamental way and worthy of special attention.

In particular, Type Ia supernovae occur in elliptical galaxies, whereas Types II, Ib, and Ic do not. Elliptical galaxies have converted essentially all their gas into stars long ago and to a great extent have ceased the making of stars. Thus elliptical galaxies are thought to consist of only old, low-mass, long-lived stars. The high-mass stars born long ago should be long dead. This has given rise, in turn, to the idea that Type Ia supernovae must come somehow from low-mass stars. Because spiral galaxies contain a mix of high-mass and low-mass stars, that spirals produce both Type Ia and Type II supernovae is not surprising.

Another aspect that has driven thinking about Type Ia supernovae is that their observed properties are remarkably uniform. Type Ia events tend to follow the same light curve. In addition, as Type Ia brighten and decline, the alterations in their spectra follow a very predictable course. Because white dwarfs of the Chandrasekhar mass would be essentially identical and hence undergo nearly identical explosions, the observed homogeneity of Type Ia has pointed to an origin in exploding white dwarfs. We now know that all Type Ia supernovae are not exactly identical. The reasons for this are the subject of active current research, as will be discussed later.

The most popular notion for how to turn a low-mass star into a supernova is thus to rejuvenate a white dwarf. The idea is that the more massive star in an orbiting pair could evolve and form a white dwarf. The low-mass companion could then take a long time to evolve, but the companion would eventually swell up as a red giant and dump mass onto the white dwarf. If the total mass accumulated by the white dwarf approaches the Chandrasekhar mass of about 1.4 solar masses, the white dwarf might then explode. A variation on this theme is that the white dwarf could grow in mass in a cataclysmic variable system where the mass flows from a main sequence star (Chapter 5). This process is slow, and the system could still last a long time before exploding. Yet another possibility is that Type Ia supernovae arise from systems of two white dwarfs that slowly merge due to the emission of gravitational waves generated by their orbital dance (Chapter 5, Section 4).

Careful studies of the observed properties of Type Ia supernovae are completely consistent with the general picture that the explosion occurs in a white dwarf. Near peak light, the spectra of Type Ia supernovae show elements such as oxygen, magnesium, silicon, sulfur, and calcium. These are just the elements expected if a mixture of carbon and oxygen burns to

produce somewhat heavier elements consisting of differing numbers of “helium nuclei.” As a Type Ia supernova evolves, the spectrum becomes dominated by iron and other similarly heavy elements. These elements can be produced by burning carbon and oxygen all the way to iron. The nuclear binding energy of iron is at the bottom of the “nuclear valley,” where the neutrons and protons in the nucleus are most compressed (Chapter 2, Section 4).

In the process of expanding and thinning out, the outer, more tenuous portions of a supernova are seen first, and the inner, denser, more opaque portions are only seen later. The information revealed by the evolution of the spectra is then consistent with a configuration in which the denser inner portions of the exploding star burn all the way to iron and iron-like elements, but the outer parts are composed of matter that results from carbon burning, but that is not so thoroughly processed. Computer models of exploding white dwarfs give results that match this pattern rather well. The exact nature of the combustion is still being explored, but the most successful models adopt a progenitor that is a carbon/oxygen white dwarf with a mass very near to, but less than, the Chandrasekhar mass.

At this point, I must correct a long-standing and erroneous view of the nature of Type Ia supernovae. This view is shared by many wise experts and neophytes alike because they have not followed this research closely. A casual view that permeated the astronomical community and the popular astronomical literature decades ago, and that is very difficult to root out, is that to make a Type Ia supernova, matter is added to a white dwarf until the Chandrasekhar mass is exceeded and the white dwarf collapses. *This is wrong!* The reason this notion is so persistent, I suspect, is that the idea of exceeding the mass limit and collapsing is simple and visceral. In addition, the “other” means of making supernovae does involve core collapse, and so it is easy to confuse the two mechanisms. There are also circumstances where some white dwarfs might collapse, but if so, the process does not yield the events we observe as Type Ia supernovae. Rather, mass is added, we believe, increasing the density in the center of the white dwarf until finally carbon can ignite. This condition of carbon ignition and subsequent unregulated thermonuclear runaway happens when the white dwarf has a mass about 1 percent less, not more, than the Chandrasekhar mass, and it blows the white dwarf up completely, so there is no collapse. This is a somewhat more complicated and perhaps less intuitive process (think dynamite!), and this may be why it has not permeated all corners of the community of interested people. Nevertheless, the supernova community stopped talking about exceeding the Chandrasekhar limit and collapse in the 1960s, and it is rather dismaying to find experts in related areas, never mind popular astronomy enthusiasts, still referring to this outmoded physical picture. The overwhelming observational evidence is that Type Ia supernovae arise from carbon/oxygen white dwarfs of mass a little less than the Chandrasekhar limit that do not collapse, but blow up completely by a process of thermonuclear explosion.

Type Ia supernovae explode because the white dwarf is supported by the quantum pressure, and any burning under those circumstances is unregulated, as we discussed in Chapters 2 and 5. For Type Ia supernovae, burning is unregulated in the extreme. As a white dwarf approaches the Chandrasekhar limiting mass, the central density gets very high. Formally, the density would go to infinity just at the Chandrasekhar limit, but in practice other physics, in this case carbon burning, will come into play. The high density triggers the ignition of carbon but also ensures that, under these circumstances, the quantum pressure will be exceedingly large. The white dwarf will have a finite temperature that will help to promote the carbon burning, but

the thermal pressure is negligible. The story of unregulated burning we have told before will then play out in the most drastic way. The carbon begins to burn and to release energy. The quantum pressure does not budge. There is no mechanical response to expand and cool the star and damp the burning. The burning goes even faster, raising the temperature even more and producing ever faster burning. Under the extreme conditions at the center of a white dwarf with a little less than the Chandrasekhar mass, the burning cannot be controlled, the oxygen also ignites, and all the fuel is consumed to iron-peak elements in a flash. The result is a violent thermonuclear explosion.

There are two different ways of propagating a thermonuclear explosion in a white dwarf. One is a subsonic burning like a flame, a process called a *deflagration*. The other is a supersonic burning that is preceded by a shock front, very much like a stick of dynamite. This process is known as a *detonation*. We have known since the 1970's that Type Ia explosions cannot be the result of pure detonation. The supersonic burning rips through the model white dwarf before it can expand and adjust and essentially the whole star is converted to iron-like matter. That is not what we see! We must account for the oxygen, silicon, sulfur, and calcium in the outer layers. The most sophisticated current models, those that best match the data, have the unregulated carbon burning begin as a boiling, turbulent deflagration and then make a transition to a supersonic detonation, as illustrated in Figure 6.5. These are known as *deflagration-to-detonation models*.

Both deflagration and deflagration-to-detonation models naturally create iron-like matter in the center and intermediate elements like magnesium, silicon, sulfur, and calcium on the outside. These models also predict that the white dwarf is completely destroyed, leaving no compact remnant like a neutron star or a black hole. This comparison of theory and observation thus strongly points to an interpretation of Type Ia supernovae as the explosion of a carbon/oxygen white dwarf at just less than the Chandrasekhar limit.

There are ways to distinguish white dwarfs that explode only by subsonic deflagration and those that explode in the more complex deflagration-to-detonation picture. The deflagration pushes matter out ahead of it at nearly the speed of sound, but the burning proceeds at intrinsically less than the speed of sound so it cannot catch up with, and burn, all the expanding matter. This means that models that rely purely on deflagration to explode the supernova must leave some unburned matter, still composed of carbon and oxygen, in the outer, fast-moving layers. The deflagration models also tend to leave “fingers” of unburned carbon and oxygen extending down to the center of the explosion. My colleagues at the Naval Research Laboratory, Vadim Gamezo and Elaine Oran working with Alexei Khokhlov, have shown that deflagration-to-detonation models drive a detonation through the “fingers” of unburned matter left by the deflagration phase and through the outer layers. The result is to scour the unburned carbon from the ejected matter. Observations in the infrared are a powerful way to look for carbon. Observations and analysis by my colleagues here in Texas, Howie Marion and Peter Höflich, have shown that carbon seems to exist neither in central “fingers” nor in the outer layers of normal Type Ia supernovae. Some outer, high velocity carbon is seen in some “sub-luminous” events, but this is naturally accounted for in deflagration-to-detonation models by triggering the detonation somewhat later. At this writing the evidence seems to strongly favor some version of the deflagration-to-detonation models. The physics of when and why the explosion makes the transition from deflagration to detonation remains to be solved satisfactorily.

Convergence on deflagration-to-detonation models for the explosion of does not, however, answer all the mysteries about the nature of Type Ia supernovae. For Type II supernovae, we think we understand the broad outlines of the evolution of massive stars to form collapsing iron cores. We do not understand how the collapsing core results in an explosion. For Type Ia supernovae, the situation is just the opposite. There is nearly unanimous agreement that the mechanism of Type Ia supernovae is a violent thermonuclear explosion that obliterates the star. Despite this convergence of opinion on the mechanism, there is no generally accepted picture of the evolutionary origin of these peculiar events. The question of how the white dwarfs grow to the Chandrasekhar mass is still a knotty, unsolved problem. There has been no direct evidence that Type Ia supernovae arise in binary systems. Despite this lack of direct evidence, all the circumstantial evidence points to evolution in double-star systems, and there are few credible ways of making a white dwarf explode without invoking a binary companion. The challenge is to figure out what binary evolution leads to a Type Ia explosion.

New perspectives on the nature of Type Ia supernovae came with evidence produced in the 1990s that confirmed a long-standing suspicion. Type Ia supernovae are not all identical. They show interesting variations that are mostly subtle, but real. In some cases, the variations are not even so subtle. The general trend is that Type Ia supernovae that are brighter than average decline from maximum brightness a bit slower than average. The events that are a bit dimmer than average (some by as much as a factor of 2) decline more rapidly. Models of exploding Chandrasekhar-mass white dwarfs can account for this behavior if the explosion in some stars makes the transition from a subsonic deflagration to a supersonic detonation a little earlier than in others. Why this should be so is the object of current research.

The observed variety of Type Ia behavior seems to correlate with the nature of the host galaxy. Elliptical galaxies seem to produce selectively Type Ia supernovae that are of the dimmer, more rapidly declining variety. Within spiral galaxies, the inner portions seem to produce the full range of behavior, but the outer parts of the galaxy produce especially homogenous explosions. We do not yet understand all the variables, but there is probably a variety of ways of making white dwarfs explode, and the progenitor systems can display a range in ages. Some Type Ia supernovae may come from mass transfer in “normal” binary systems, from some variation on a cataclysmic variable. Others may come from merging white dwarfs. Some may come from stars near 8 solar masses that have relatively short lifetimes and others may come from stars with closer to 1 solar mass that have lifetimes approaching that of the Universe itself.

The task of figuring out the prior evolution of Type Ia supernovae is made harder if one accepts that the supernovae arise in white dwarfs of the Chandrasekhar mass. Recall from Chapter 5 that the average white dwarf has a mass of only 0.6 solar masses. This means that the mass must more than double if the process starts with one of these white dwarfs. The task might be made easier if the white dwarfs born in binary systems are systematically more massive. There is some evidence that this may be the case. Note that if the white dwarf is in a system that undergoes a nova explosion every 10,000 years or so, the mass of the white dwarf could actually decrease! This is not an easy problem.

For this reason, there has been considerable attention paid to mechanisms that would lead a white dwarf to explode even though it had less than a Chandrasekhar mass. The most likely such model is one where a white dwarf accretes mass rapidly enough that the accreted hydrogen

remains hot and supported by its own thermal pressure. The hydrogen then burns on the surface of the white dwarf in a regulated manner, and a nova explosion is avoided. Under these circumstances, however, a thick layer of helium can build up surrounding the inner carbon/oxygen core. The helium layer can be supported by the quantum pressure. If this helium ignites, computer models show that a violent explosion occurs. The explosion not only burns the helium but can send a shock wave inward that causes the inner carbon/oxygen white-dwarf core to burn as well. All this happens very quickly, a matter of seconds, so the result is a single powerful explosion. This is a very plausible mechanism to produce an explosion. The problem is that this mechanism does not produce results that are in good agreement with the observations. The helium burns to iron-like material on the outside that should be seen first and produces only thin layers of intermediate elements like silicon and calcium that are ejected with the wrong velocities. The ejecta tend to be too hot as well. Despite the appeal of these models, nature seems to prefer exploding white dwarfs of nearly the Chandrasekhar mass.

There are currently two “best bets” for how to generate Type Ia supernovae. Both involve mass transfer onto a white dwarf in a binary system. One invokes transfer of hydrogen from a red giant at just the right rate. The mass transfer must be rapid enough that the collected hydrogen does not undergo a nova explosion that ejects the hydrogen along with part of the white dwarf. Apparently, the mass transfer must be rapid enough that even the helium remains hot, supported by the thermal not the quantum pressure, so that igniting the helium does not cause an explosion with the wrong properties. If the mass transfer is too rapid, however, a common envelope of hydrogen will engulf the white dwarf. The hydrogen should show up in the explosion. That would be a violation of the basic observational definition of a Type I supernova. There may be binary configurations where the mass transfer is “just right.” The hydrogen will burn gently to helium, the helium will burn gently to carbon and oxygen, and that carbon and oxygen will settle onto the core to cause the core to grow toward the Chandrasekhar mass. Candidate systems have even been identified among a special class of X-ray sources called *supersoft X-ray sources*.

An interesting clue to this problem was provided by the discovery by Mario Hamuy of Carnegie Observatories of a supernova that had obvious evidence for hydrogen, but, when one looked, an underlying spectrum that was that of a Type Ia. The hydrogen was apparently transferred from an ordinary red-giant companion. Polarization observations by our group showed that the hydrogen was distributed in an extended disk as one might think appropriate for a strong mass transfer that slopped matter out of the binary system as well as onto the white dwarf. It is not clear how common this sort of explosion is, although a few other candidates have been identified. It is also true that while the hydrogen was totally obvious in this event, careful searches for wisps of hydrogen have failed to produce any evidence in normal Type Ia.

Another line evidence concerning the binary nature of Type Ia has been found by the recent discovery of high velocity shells containing calcium that are somehow detached from the supernova ejecta. Where data has been obtained, these shells show polarization and hence some break-down in spherical symmetry. My Texas colleagues Chris Gerardy (now at University College London) and Peter Höflich have argued that the calcium is in a shell otherwise composed of hydrogen (or perhaps helium) that pre-existed in the binary system and was compacted and ejected by the supernova explosion. In models, the calcium radiates efficiently in the compacted shell and the hydrogen (or helium) radiates more feebly and remains invisible. This high-velocity calcium thus may be a clue to the nature of the binary system and hints that

the system contains a hydrogen-rich star, even though the hydrogen is not directly detected. The swept up matter revealed by its calcium emission may come from an accretion disk, from the companion star, or perhaps from matter that was previously part of a common envelope that still lingers nearby.

The other popular model for producing a Type Ia supernova is by the merging of two white dwarfs in a binary system (Chapter 5, Section 4). This merging must happen sometimes. Some binary white dwarfs are seen. There is still controversy concerning whether there are enough binary white-dwarf systems with total mass exceeding the Chandrasekhar mass to produce Type Ia supernovae at the observed rate. In addition, the process by which the smaller-mass white dwarf fills its Roche lobe and comes apart, dumping its mass on the larger-mass white dwarf as described in Chapter 5, is complex and not well understood. The disrupted matter will swirl around the larger-mass white dwarf in a thick disk. How that matter will settle onto the remaining white dwarf is not completely clear.

Yet another way to pursue evidence that Type Ia explode in binary systems is to look for the left-overs; not a compact remnant, but the companion star that would be left behind if the explosion occurs in a mass-transferring binary system. The matter in stars is rather concentrated toward their centers and that makes them tough. A nearby supernova could strip off some matter from the outside, but a companion star will easily survive the explosion. On the other hand, the companion star will be released from its orbit when the binding gravity of its companion disappears in the explosion. The companion should thus be slung out of the site of the explosion. The companion might be a rather normal little red main sequence star as observed in many cataclysmic variable systems, but there are billions of them in the Galaxy so identifying the companion is not a simple thing to do. Pilar Ruiz-Lapuente from the University of Barcelona and her colleagues focused on the remnant of Tycho's supernova. Using images from the Hubble Space Telescope, they did not find any red giants that could be the companion. but they did identify a yellow star much like our Sun that is moving out of the vicinity of the explosion at about three times the average speed of other nearby stars. They suggest that this star is the surviving companion of Tycho's supernova.

The accumulating clues thus suggest that Type Ia do arise in binary systems and that the most common configuration involves mass transfer from a relatively normal companion star. White dwarf mergers might contribute to some small fraction of Type Ia explosions, but there is no firm evidence for that at this time.

7. Light Curves: Radioactive Nickel

Supernovae display a variety of shapes to their light curves. Type Ia supernovae are the brightest. They decay fairly rapidly in the first 2 weeks after peak light and then more slowly for months. Some Type II supernovae have an extended plateau and some drop rather quickly from maximum light. Both types seem to have a very slow decay at very late times, several months after the explosion. Types Ib and Ic supernovae are typically fainter than Type Ia by about a factor of 2, but they have similar shapes near peak light and show evidence for a slow decay at later times. These patterns tell us something about the star that exploded and about a fundamental process that is probably taking place in all of them: radioactive decay.

When a supernova first explodes, the matter is compact, dense, and opaque. To reach

maximum brightness, the ejected matter must expand until the material becomes more tenuous and semitransparent. The size the ejecta must reach is typically 10,000 times the size of the Sun. This is 100 times the size of a red giant and 100 times the size of the Earth's orbit. As the matter expands, however, it cools. If the matter must expand too far before heat leaks out as radiation, the material may have cooled off so that there is no more heat to radiate.

Most Type II supernova explosions are thought to occur in red-giant envelopes. These are very large structures. After the explosion, large envelopes do not have very far to expand before they become sufficiently transparent to leak their heat as light. As they begin to radiate, Type II supernovae still retain a large proportion of the heat that was deposited by the shock wave that accompanied the supernova. Near maximum light and on the typical plateau that lasts for months, Type II supernovae shine by the shock energy originally deposited in the star. The deposited energy presumably arises in the core collapse process.

For a Type I supernova, however, the story is different. Whether the exploding star is a white dwarf, as suspected for a Type Ia, or the bare core of a more massive star, as suspected for Types Ib and Ic, the exploding object is very small. The expected sizes range from one-tenth to one-thousandth of the size of the Sun. These bare cores are vastly smaller than the size to which they must expand before they can leak their shock energy. The result is that the expansion strongly cools the ejected matter, and by the time the matter reaches the point where it could radiate the heat, the heat from the original shock is all gone. This kind of supernova requires another source of heat to shine at all. All the light from Type I supernovae comes from radioactive decay.

The nature of a thermonuclear explosion is to burn very rapidly. If the explosion starts with a fuel built from multiples of helium nuclei – carbon, oxygen, or silicon – that has equal numbers of protons and neutrons, then the immediate product of the burning will also have equal numbers of protons and neutrons. This is because the rapid burning takes place on the time scale of the strong nuclear reactions. To change the ratio of protons to neutrons requires the weak force and thus a longer time. Nature, however, does not leave the burned matter with equal numbers of protons and neutrons. Rather, nature prefers to form the element with the most tightly compacted nucleus, that of iron, which has twenty-six protons and thirty neutrons.

Nature manages to make iron in a thermonuclear explosion in a three-step process. The first step is to forge an element that is close to iron but that has equal numbers of protons and neutrons. This element, like iron, has a nucleus that is tightly bound by the nuclear force and has the same total number of protons plus neutrons, fifty-six, but with twenty-eight protons and twenty-eight neutrons. This is the element that will form first, before the slower weak interactions come into play. This condition singles out one element, nickel-56. The unregulated burning of carbon or oxygen or silicon will naturally first produce nickel-56.

Nickel-56 is, however, unstable and therefore undergoes radioactive decay. The radioactive decay is induced by the weak force. One of the protons in the nickel converts to a neutron. The result is the formation of the element cobalt-56 with $28 - 1 = 27$ protons and $28 + 1 = 29$ neutrons. In the process, an electron is absorbed to conserve charge, and a neutrino is given off to balance the number of leptons. Excess energy comes off as gamma rays, high-energy photons. The gamma rays can be stopped by collision with the matter being ejected from the supernova and their energy used to heat the matter. The hot matter shines as the light we observe

on Earth. The power of the light falls off as the nickel decays away and as the matter expands so that it is less efficient in trapping the gamma rays. The neutrino always just leaves the star and plays no role in this heating.

The cobalt-56 that forms is also unstable. Again, the weak force induces a proton to convert to a neutron. The result has $27 - 1 = 26$ protons and $29 + 1 = 30$ neutrons. This is just good old iron-56, nature's ultimate end point. This decay again produces a neutrino and gamma-ray energy. In this case, charge is conserved by emitting an antielectron, or positron. The positron will quickly collide with one of the electrons that are floating around normally, one for every proton. The annihilation of the electron will produce another source of gamma rays. Iron-56, with twenty-six protons and thirty neutrons, sits at the bottom of the nuclear energy valley, and so it is stable. This radioactive decay scheme, nickel to cobalt to iron, is just one of nature's ways of rolling things down the nuclear hillside to become iron.

The radioactive decay of these elements is controlled by a quantum uncertainty. One does not know what atom will decay, but on the average half will decay in a given time. For nickel-56, the time for half to decay is 6.1 days. After another interval of 6.1 days, half of the remaining half will decay so that after 12.2 days only one-quarter of the original nickel remains. After 18.3 days, only one-eighth of the original nickel will survive. This time scale, about a week, is the time for the gamma rays from the radioactive decay to pump energy into the exploding matter. Likewise, the cobalt-56 decays with a half-life of about 77 days, roughly 2 months. These times are long compared with the times for the basic explosion to ensue, a matter of seconds. That is why the nickel-56 forms first in this type of explosion and the iron forms only later, over several months. The observed light curves of Type I supernovae decay somewhat faster than the decay of nickel-56 in the early phase and of cobalt-56 in the later phases. The reason is that not all the gamma rays produced in the decay are trapped and converted to heat and light. Some of the gamma rays escape directly into space.

For Types Ib and Ic, the amount of nickel required to power the light curve is about one-tenth of the mass of the Sun. This amount of nickel is consistent with many computations of iron-core collapse. The nickel is produced when the shock wave, of whatever origin, impacts the layer of silicon surrounding the iron core. Type Ia supernovae are generally brighter and must produce more nickel, of order 0.5–1 solar mass. The dimmest Type Ia events require only 0.1–0.2 solar mass of nickel. The models of Type Ia supernovae based on thermonuclear explosions in carbon/oxygen white dwarfs of the Chandrasekhar mass produce this amount of nickel rather naturally in the explosion. The amount can vary depending on, for instance, the density at which the explosion makes the transition from a deflagration to a detonation, so the variety of ejected nickel mass can also be understood, at least at a rudimentary level.

If Types Ib and Ic are related to the cores of massive stars, as the circumstantial evidence dictates, then their explosion mechanism should be similar to that of Type II supernovae. This suggests that Type II should also eject about 0.1 solar mass of nickel-56. This is not enough to compete with the heat and light from the shock near maximum light, but as the ejected matter continues to expand and cool, the shock energy dissipates, and the supernova gets dimmer. At this phase, the dimmer but steady source of radioactive decay should take over. The evidence from fading Type II supernovae shows that this is the case. Once again, not all the gamma rays are trapped. Some must radiate directly into space. A properly designed gamma-ray detector flown in orbit should see these missing gamma rays and directly confirm the validity of this

picture. As we will see in Chapter 7, this was the case for SN 1987A.

WHEN BETELGEUSE BLOWS

[rffFor years, every time I gave a popular lecture on supernovae, someone would ask, “What will happen to the Earth when a nearby supernovae explodes.” Each time I would say, “I thought about that a little a long time ago, but I really need to work that out, so I know how to answer this question.” Then after the lecture, I would return to work-a-day issues and forget until the next popular lecture. To get a record down on paper that I can use in the next lecture, here is a sketch of what will happen when the most likely nearby star explodes.

Betelgeuse is a red-giant star that marks the upper-leftmost shoulder of the constellation of Orion as we look at it from Earth. You can see it easily from anywhere in the northern hemisphere on a winter or spring evening. We do not know the precise mass of Betelgeuse, but we can make an intelligent guess. That will give us a good guess as to its fate and what will happen at the Earth.

Thanks to careful measurement by triangulation we know quite accurately how far away Betelgeuse is. It is 427 light years away. That is long by human standards, but right next door in a Galaxy that is 100,000 light years across. There are closer stars, but none that are likely to explode. At this distance, Betelgeuse presents little threat to the Earth, but we will sure notice it when it goes off. It is a good example of the low-level impact that will contribute to the stochastic history of bombardment of the Solar System by astronomical events over its 5-billion-year history. Such events should occur roughly once every million years.

From the power received at Earth over all wavelength bands and its distance, we can estimate that Betelgeuse emits a luminosity of about 50,000 to 100,000 times that of the Sun. From computer models, we can further estimate that this luminosity in a red giant requires a star of original main sequence mass of about 15–20 solar masses. This mass is such that, in the absence of a stellar companion, and Betelgeuse seems to have none, there will be little mass loss to winds, so this is probably a pretty good estimate. Stars in this mass range are predicted to evolve iron cores and undergo core collapse to form a neutron star and an explosion. Betelgeuse is nearly a canonical candidate for a Type II supernova explosion. We do not know exactly when it will explode. The final stages after a star of this mass becomes an extended red giant are typically no more than 10,000 years. We do not know when in the next 10,000 years it will explode (it may be tomorrow!), but we can estimate the progression of events when it does.

Upon core collapse, Betelgeuse will emit 10^{53} ergs of neutrinos, each with an energy characteristic of a nuclear reaction. This burst of neutrinos will take about an hour to pass through the hydrogen envelope and into space. They will arrive in the Solar System 427 years later and be the first indication that Betelgeuse has erupted. These neutrinos will deliver about 2×10^8 recoils in the body of a 100-pound woman. This effective level of radiation exposure is far less than a lethal dose (by a factor in excess of 1,000, depending on how the energy is actually deposited) but might cause some chromosomal damage. The shock wave generated by the collapsing core and the formation of a neutron star will require about a day to reach the surface. The breakout of that shock will generate a flash of ultraviolet light for about an hour that will be about 100 billion times brighter than the total luminosity of the Sun. This burst may not exceed the ultraviolet light from the Sun at the Earth, but could affect life on outer satellites, if there is

any, or any explorers from Earth, if we have ventured far from the Sun by the time this happens. This blast of ultraviolet might cause some disruption of atmospheric chemistry. The ejecta of the supernova will expand and cool after shock breakout, and the total luminosity will first dim and then rise to maximum in about 2 weeks as the supernova material expands to about 100 times the Earth's orbit and the photon diffusion time through the expanding matter becomes comparable to the time required for appreciable expansion of the matter. The total luminosity will then be about a billion times that of the Sun. At its distance, Betelgeuse will be a factor of about one million dimmer than the Sun, magnitude – 12, about the same as a quarter Moon. This phase will last during the “plateau” phase of the light curve, 2 or 3 months. The observed surface of the supernova during this interval will be roughly constant at an effective temperature of about 6,000 K, slightly hotter than the Sun. After the hydrogen envelope has expanded and electrons and protons have all recombined to make neutral hydrogen atoms, the envelope will be nearly transparent, and the light curve will begin a rapid decline.

In a typical supernova of this type, the emission is dominated for the next year or so by radioactive decay of cobalt to iron (nickel will have already decayed away). The expanding envelope of hydrogen is likely to remain opaque to these gamma rays until substantial decay has occurred, so such an event is unlikely to provide a substantial source of gamma rays. If Betelgeuse produces a bright pulsar (Chapter 8), it might be a substantial source of gamma rays for thousands of years.

The ejecta from Betelgeuse will freely expand for about 1,000 years and span about 20 light years in that time. During this time, the ejecta will be cold and dim. The supernova material will then start to pile up appreciable mass in interstellar matter and enter the supernova remnant phase. The supernova remnant will turn on as an X-ray source and begin to produce cosmic rays by acceleration of particles at the shock front. The supernova material will slow down, but a shock will race ahead into the interstellar matter, decelerating as it sweeps up ever more mass. The shock wave in the interstellar matter will be fully developed in about 20,000 years when it has expanded to about 30 light years. The shocked matter will begin to radiate substantially and cool off when it has expanded to about 100 light years about 100,000 years after the explosion. The remnant will plow on through the interstellar matter. The shock from Betelgeuse will be very mild by the time it reaches the Solar System and will probably be easily deflected by the solar wind and magnetopause. The exception might be if there is a low-density, interstellar “tunnel” between us and Betelgeuse that would channel some of the energetic matter to us before it slowed down.

All these effects would be much stronger if the supernova were only 30 light years from the Earth. There are no candidate stars around us now, but on its galactic journey, such nearby explosions have probably happened several times in the 5 billion year life of the Earth. Such events could be dangerous by triggering harmful mutations, but they might also be helpful because evolutionary “shocks” can also single out healthy mutations and drive biocomplexity. The Earth is coupled to this complex galactic environment, and the story of life on Earth will not be fully known until such long-term, sporadic effects are understood.