8

Neutron Stars

Atoms with Attitude

1. History – Theory Leads, for Once

In 1932, the brilliant Russian physicist Lev Landau argued on general grounds that the newly discovered quantum pressure could not support a mass much in excess of 1 solar mass. He addressed his discussion to electrons, but the type of particle did not matter. In 1933, the neutron was discovered, after Landau's paper had been submitted. In retrospect, Landau's arguments applied to the quantum pressure of neutrons as well. An object supported by the quantum pressure of neutrons should be smaller and denser than a white dwarf, but it should have nearly the same maximum mass, about 1 solar mass.

Fritz Zwicky of Caltech was one of the world's first active supernova observers. Quick on the pick-up, Zwicky suggested in 1934 that supernovae result from the energy liberated in forming a neutron star. Not until a year later, in 1935, did the precocious young Indian physicist, Subramanyan Chandrasekhar, present his rigorous derivation of the nature of the quantum pressure and the mass limit to white dwarfs that bears his name.

Robert Oppenheimer made history with his leadership of the Manhattan Project, but among his most widely known papers are two published with students in 1939. The first of these papers used the complete theory of general relativity for the first time to estimate the upper mass limit of neutron stars to be 0.7 solar mass. The second paper explored the result of violating that limit with the resulting production of a black hole. The upper limit to the neutron star is now commonly referred to as the Oppenheimer-Volkoff limit, after the authors. In the 1960s, repulsive nuclear forces between the neutrons were added to the purely quantum effects. As a result, the estimates of the maximum mass of neutron stars rose to between 1.5 and 2.5 solar masses.

In 1964, John Archibald Wheeler suggested that the power radiated by the Crab nebula could plausibly be provided by the rate of loss of rotational energy of a neutron star. This proved to be a prescient guess. At about the same time, Rudolph Minkowski, an old cohort of Fritz Zwicky, was studying the Crab nebula. He pointed out that, although most of the stars seen in a photograph were foreground or background stars, one, apparently buried in the heart of the nebula, had a peculiar spectrum and an abnormally blue color. Minkowski could not prove that this peculiar star was in the nebula. There was not a shred of rational evidence relating Wheeler's speculation to Minkowski's observations, but the relation turned out to be true.

Theoretical astrophysicists often find themselves dragging along behind the observations, trying to explain some exciting new phenomenon ex post facto (quasars represent a superb example). In the case of neutron stars, however, the theorists were way out in front. More than three decades passed from the first theoretical discussions of neutron stars until some confirming evidence came in.

In 1967, Jocelyn Bell was a graduate student working with Anthony Hewish on a peculiar radio telescope at the University of Cambridge in England. The telescope was a series of wires run helter skelter, designed to look for rapid modulation of radio signals by the solar wind. What Ms. Bell noticed among the reams of data was a source of regularly pulsed radio emission. The pulses lasted 0.016 seconds and recurred quite regularly, every 1.33730115 seconds, with astounding accuracy.

The investigators were mystified at first and then, after some contemplation, petrified. There had been a long-standing expectation that any extraterrestrial civilization would signal its existence with some regularly modulated mechanism. The strange signals were dubbed LGMs, short for little green men, and a strong air of secrecy cloaked the lab. This conclusion was too significant to be blabbed about, while further checks ensued.

Soon other such sources were discovered. Significantly, and much to the relief of the researchers, they found the pulse periods were gradually increasing. The fantastically accurate period was not locked in as it would be with an artificial mechanism, but slowly drifted. Whatever these things were, they represented a natural phenomenon. The discovery of *pulsars*, pulsating radio sources, was announced to the world. Anthony Hewish won the Nobel Prize for Physics for the discovery of neutron stars as pulsars in 1974. To the discomfit of some, Jocelyn Bell, whose perspicacity revealed the unexpected signal, did not share in the award. Dr. Bell, a gracious woman, went on to a fruitful career as an X-ray astronomer.

2. The Nature of Pulsars – Not Little Green Men

What were these pulsars? They could not be ordinary stars. Even the light travel time across the Sun is a few seconds, and the pulses in these objects lasted only a fraction of a second. More practically, the fastest motion the Sun could withstand would be if it changed substantially in about a half hour. This is the Sun's dynamical time scale, the time it requires to respond to an imbalance between gravity and pressure. Any global motion of the whole Sun on a faster time scale, whether by rotation, oscillation, or any other mechanism, would mean that the Sun would tear apart.

White dwarfs are more compact and able to withstand rapid movement. One second – a characteristic time between pulsar pulses – is just about the natural time scale for a white dwarf. Just after the discovery of pulsars there was a great flurry of activity exploring white dwarf models for pulsars. The white dwarfs were pictured to be rotating or oscillating. Some people even considered neutron stars. Because neutron stars were even more compact, they would have no trouble responding quickly enough. The natural dynamical time scale for an oscillating neutron star is about 1 millisecond, or 0.001 second, so there was some question why a neutron star should respond as slowly as 1 second. At first, neutron stars were considered a radical, though not impossible, explanation for pulsars.

The studies that showed that the periods of pulsars lengthened with time continued as the theorists thrashed around for a consistent explanation of pulsars. The gradual lengthening of the time between pulses turned out to be a key, if subtle, clue. Studies of oscillating stars show that they tend to respond more rapidly as they lose energy. The reason is that the oscillations themselves tend to make the star somewhat more bloated and unresponsive. As the oscillations die away, the star gets more compact and bounces more quickly. A rough analogy is to drop a ball and listen for the bounces; they become closer together as the ball bounces less and less high. The lengthening of time between pulses suggested that the pulsar phenomenon had nothing to do with oscillations. As a rotating object loses energy, it spins more slowly, and so the time to make one revolution lengthens. This is in accord with the behavior of pulsars, so some rotational phenomenon was considered the most likely explanation for pulsars.

The next major breakthrough came from studies of the Crab nebula. Ten or twenty pulsars had been discovered, all with periods of about 1 second. Then astronomers focused on the strange star Minkowski had pointed out years before. The star turned out to be a pulsar! The period of the pulses was much faster than had been seen in any other pulsar. The time between pulses was only 0.033 seconds. This time is so short that no white dwarf could oscillate or rotate that fast without being torn apart. The pulsar in the Crab nebula had to be a neutron star, and so, presumably, did all the others! Only rotating neutron stars could account for the whole range in periods from fast to slow. A big star cannot rotate rapidly, but a compact star like a neutron star can rotate rapidly or slowly, depending on circumstances.

The pulsar in the Crab nebula rotates relatively rapidly because it was born only a short time ago and has not had time to lose much rotational energy. The pulsars with spin periods of about a second are deduced to be 1 million to 10 million years old. The Crab pulsar is so energetic that it emits pulses of optical light as well as radio radiation.

We still do not understand clearly why the radiation comes from the pulsars in pulses. That radiation comes from the pulsars at all is, however, a clue to another important property. The neutron stars must contain strong magnetic fields to generate radiation. Fundamentally, radiation is caused by wiggling a magnetic field. This causes a wiggling electric field, which in turn causes a wiggling magnetic field, which causes a . . . Coupled wiggling electric and magnetic fields are at the heart of the process of electromagnetic radiation. Without a magnetic field, the rotating neutron star could not emit the kind of radio radiation observed. Thus pulsars must be *rotating, magnetized neutron stars*. That the pulsars are magnetic is not too surprising. Ordinary stars like the Sun generate magnetic fields. If such a star were compressed to the size of a neutron star, the magnetic field would be amplified by a factor of about 10 billion. The resulting magnetic field would be just about what is required to generate the radiation in pulsars. Whether squeezing the field of the star that collapsed to form it is the origin of the magnetic fields of pulsars is still not clear. The newly born neutron stars may act like dynamos and make their own magnetic fields.

The simplest magnetic field a neutron star could have is a so-called dipole field like a bar magnet, with a north pole and a south pole, as shown in Figure 8.1. The lines of magnetic force for such a field are arching loops, out one pole and into the other, exactly like the pattern of iron filings around a bar magnet. If the magnetic field is perfectly aligned with the axis of rotation, there will be no radiation, at least no pulsed radiation. The reason is that the magnetic configuration is too symmetric. If the magnetic field is perfectly aligned, there is no effective change in the magnetic field as the neutron star rotates. A wiggling magnetic field is required to generate radiation, and a perfectly aligned magnetic field causes no wiggles as the neutron star rotates.

Radiation will occur if the axis of the magnetic field is tipped with respect to the rotation axis. Then as the neutron star rotates, the magnetic field points in different directions, and the magnetic force at any given point in space varies continuously. This misalignment is not so special a requirement when one considers that the magnetic poles of the Earth are not lined up exactly with the rotation axis and that the magnetic poles even occasionally swap ends.

If pulsar radiation comes from the magnetic poles, we can even understand the pulses because the magnetic poles sweep around like beams from a lighthouse. A pulse would be detected every time a radio “light house beacon” pointed at the Earth. This is the most popular view of the origin of the pulses. Theories have been constructed in which the rotating magnetic fields generate huge electrical fields right at the magnetic poles. The energy in the electric field is so great that it can rip electrons from the neutron star surface or create electron/positron pairs. The particles cause a gigantic spark as they flow along the electric and magnetic fields toward the neutron star or out into space. The spark, like a bolt of lightning at the pole, emits a burst of radio static. This is the particular mechanism envisaged by which the magnetic field “wiggles” and gives rise to radiation.

There is still debate as to exactly where and how this spark forms. As the pulsar rotates around, the magnetic lines of force are carried around with it. Any charged particles caught in the magnetic field are forced to spiral along the field, but they cannot move across the field. The result is that as the neutron star rotates, the particles must rotate as well. All the particles locked to the rotation of the neutron star make a complete circle in the same time, but to accomplish this, the more distant particles, with a greater circumference to travel, are forced to move at tremendous velocities. At not too great a distance from a neutron star, the particles would be whipped around at the speed of light. The path on which particles locked to the neutron star's rotation would move at this limiting speed is known as the *speed of light circle*. The distance would be a thousand miles in the case of the Crab pulsar and 30,000 miles – roughly the Earth's diameter – for a pulsar with a period of 1 second. Because particles cannot move at the speed of light, the particles must be ripped off the magnetic field lines at the speed of light circle. The wrenching process involved would generate radiation. Some theories argue that the great forces generate electron/positron pairs and accelerate them near the speed of light circle so that the “spark” occurs there. Other theories argue that the particles to be accelerated are those pulled from the neutron star so that the spark arises closer to the neutron star surface.

By now some 600 pulsars have been discovered. Most of these are nearby in the Galaxy because their radiation is relatively feeble and cannot be detected from great distances. Extrapolation from the known number of pulsars leads to the estimate that as much as 1 percent of the mass of the Galaxy may be in the form of neutron stars, about one billion of them all told. Most of these would be “dead” pulsars, which could no longer radiate. Pulsars live about 1 million to 10 million years before their magnetic fields decay away or become aligned with the rotation axis, so that no pulses of radiation are possible.

3. Pulsars and Supernovae – A Game of Hide and Seek

When supernovae explode, they inject a large amount of matter and energy into the surrounding gas of the interstellar medium. An explosive “cloud” plows out into the interstellar gas much like a mushroom cloud rises from a hydrogen bomb on the Earth. For a bomb on Earth, the “cloud” rises upward from the ground; for a supernova, the cloud expands outward in all directions. The resulting expanding remnant of a supernova is marked by radiation in the radio that occurs when the shock wave from the supernova compresses and heats the interstellar gas and sends electrons spiraling around the interstellar magnetic field at nearly the speed of light. Interior to the shock wave that marks the point of collision, the shocked gas is so hot it emits X-rays. A *supernova remnant* can span several light years.

These extended supernova remnants live only about 100,000 years before they fade into the general interstellar gas. Pulsars "live" for 1 million to 10 million years. After that time, the neutron star is still around, but it no longer emits radio pulses. Thus pulsars live for about ten times longer than the extended remnants. One expects most pulsars not to be associated with an extended remnant, but that every extended remnant in which a pulsar was born should still surround that pulsar. Most pulsars are not associated with extended supernova remnants, as expected. Strangely enough, the converse is also true. Most extended remnants show no sign of a pulsar. The Crab nebula is a conspicuous exception to this rule. This negative conclusion has been strongly reinforced by searches for pulsars with X-ray satellites.

This is a puzzling observation. Either no neutron stars are formed in many supernova explosions, or they are not rotating or magnetic so that they cannot emit radio pulses or related traces in the X-ray band, or the pulsars pick up such a high velocity that they escape out of the gaseous remnant. It is possible that in many cases the radio radiation from pulsars is "beamed" so that it does not shine toward the Earth. On the other hand, the X-ray radiation, similar to that emitted strongly from the Crab nebula, shines in all directions so it would be difficult to hide. This raises yet another question. If pulsars are born at the same rate as supernovae explode, but many supernovae do not make pulsars, then apparently there is a way of making pulsars without the associated explosion and optical outburst that identify a supernova. No one knows how this is accomplished, if, indeed, it must be.

This is the context in which one considers the situation with Cas A and SN 1987A. All the evidence is that Cas A represents the explosion of a star of about 20 solar masses. Such a star is predicted to make a neutron star, but until now, none has obviously been seen. The same arguments apply to SN 1987A in a somewhat different context because that supernova is still so young. SN 1987A came from a star of about 20 solar masses. It emitted neutrinos, so we know it had a gravitational collapse, yet any neutron star must be at least ten times dimmer than the 1,000-year-old pulsar in the Crab nebula. Does this mean neutron stars exist in Cas A and SN 1987A but are especially dim? Does this dimness apply to the lack of observed neutron stars in older supernova remnants? Or did Cas A or SN 1987A ultimately create a black hole, and, if so, does this apply to the older supernova remnants? These questions remain central to the study of the final evolution of massive stars.

A new chapter in this story will be written by the *Chandra X-Ray Observatory* launched on July 23, 1999. The very first image obtained by *Chandra* was of Cas A, and to everyone's 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. It is not clear whether this source is a neutron star or a black hole. The *Chandra* web site asks for readers to vote between these choices. That is an amusing exercise, perhaps, but it is not the way science is done.

This point of X-ray light will be the subject of intense investigation, but a few conclusions are immediately clear. This source is ten thousand times dimmer than the pulsar in the Crab nebula. If it is a neutron star, it is clearly not putting forth the effort to radiate that it might. Even just the heat energy stored in a newly formed neutron star could generate more light than this, never mind any pulsar radiation. On the other hand, a small rate of accretion could make either a neutron star or a black hole shine in X-rays like this, so either could be powered by the fallback of some supernova ejecta that did not quite make it. This discovery also sheds light on the situation with SN 1987A. If a compact object this dim resides in the center of SN 1987A, then it is no wonder that it has not yet been detected. Progress on the study of this point of light in Cas A will undoubtedly also help us to understand whether SN 1987A left behind a neutron star or black hole.

4. Neutron Star Structure – Iron Skin and Superfluid Guts

Neutron stars are sometimes referred to as giant atomic nuclei because, like nuclei, they are composed essentially entirely of baryons, neutrons. Because they are so massive and bound by gravity, neutron stars have a “personality” beyond that of any atomic nucleus.

Neutron stars have about as much mass as the Sun, but, because of their very high densities, they are only 10 to 20 kilometers in radius. Their very outermost layers are of nearly normal composition. There are still protons and electrons. The material is probably mostly iron because all thermonuclear processes should have gone to completion. The topmost material is probably gaseous, an atmosphere hanging above the solid surface, just as on the Earth. One major difference is that in the huge gravitational field of the neutron star, the atmosphere would be only a few meters thick. The solid surface can support mountains and other rugged terrain. Mount Everest dropped onto a neutron star surface would be crushed to a foot or so in height. Typical hills and valleys on the surface of a neutron star would range up to several inches in height.

The outer solid crust of ironlike material on a neutron star would be a few kilometers thick. An important difference in the structure of this material is that the crust is permeated by the huge magnetic field. This magnetic field alters the structure of atoms. Electrons can move along a magnetic field line but cannot move across field lines. This rule applies even to the electrons in atoms if the magnetic field is strong enough. The result is the deformation of atoms into long skinny strings, with the electron clouds elongated along the magnetic field lines and confined in transverse directions. These atoms can in turn be linked to form new kinds of long skinny molecules, which could only exist in the extreme conditions of the crust of a neutron star.

Deeper into the neutron star, electrons are squeezed tightly by the exclusion principle, and the quantum energy they acquire forces them to combine with a proton to form a neutron. The nuclear forces cannot hold a large excess of neutrons into a nucleus, so neutrons begin to leak out of specific nuclei and move around freely in the material. This process is known as *neutron drip*. The densities at which it occurs are higher than the highest density of any white dwarf, but these conditions are still found only a few kilometers deep in the neutron star.

Upon reaching depths where the density is comparable to the density of normal atomic nuclei, nothing resembling a normal atom can exist. The material is essentially all neutrons, although there is a scattering of protons and electrons. The few electrons can still exist because they are so sparsely spread that the effects of exclusion are small, and their quantum energy is not appreciable. There is one proton for every surviving electron to balance charge. The densities are so high that the exclusion effect on the neutrons is dominant and their quantum energy, moderated by effects of nuclear forces, provides the pressure to support the neutron star. The quantum uncertainty in the "cloud" that represents a massive neutron is smaller than that for the cloud of the smaller-mass electron. This is why electrons feel squeezed first, and neutrons must be raised to much higher densities before the exclusion of one neutron by another has an appreciable effect.

A remarkable transition in the nature of neutron-star material is made at higher densities. The nuclear forces between neutrons have another important role besides just altering the pressure. The nuclear forces cause the quantum waves that represent the neutrons to line up in a special way that minimizes the repulsive nuclear forces. The result is that the neutrons are thought to form what is called a *superfluid*. A superfluid is a special state of matter in which all the particles flow in consonance and the result is absolutely zero viscosity, no resistance to motion. Water has much less viscosity than molasses, but a superfluid has none at all! Physicists have created superfluids in the laboratory by cooling liquid helium to near absolute zero. This reduces the thermal energy in the helium, and helium has no interfering chemical reactions because it is a noble gas. The result is that the quantum properties dominate, and the quantum waves of the helium atoms can line up in such a way as to form a superfluid. The resulting material flows so easily that if care is not taken, it will flow up the side of the beaker and out of the experiment! Lev Landau, with whom we introduced this chapter, won the Nobel Prize in Physics for his work on liquid helium in 1962.

At the highest densities in the center of a massive neutron star, the quantum effects among the neutrons can cause yet another arrangement of the structure. Theories predict that the neutrons will clump together into a rocklike solid. This material would be somewhat akin to the solid crust. In the crust the solidification is due to electrical forces on the electrons, whereas in the core the solidification is due to nuclear forces among the neutrons. At these most extreme densities, the huge gravitational energy can be converted into mass. Exotic particles that do not normally exist in nature could spring spontaneously into existence, but there is no proof that such processes occur.

This picture of the interior of a neutron star just sketched follows from the theoretical extrapolation of known physics to extreme conditions. Fortunately, there is some evidence that the picture is at least qualitatively correct. This evidence comes from "glitches" observed in the rate of pulses from pulsars. As we have said, pulsars generally slow down with time in the sense that their pulses slowly get farther and farther apart. This effect is quite gradual, of order of one part in a million per year, and it is only due to the exceedingly accurate rate of pulses that the slowdown can even be detected. Occasionally, however, a pulsar will speed up for a short while, and the time between pulses will become shorter. After some time, the pulses will settle back into their old pattern of gradual slowing. This behavior is known as a "glitch," which means, in general, an unexpected interruption or change in behavior. Glitches have been observed in a few of the youngest pulsars. Apparently, the older pulsars have settled down into a state where they do not glitch anymore. The Crab pulsar has been observed to glitch. There is another supernova remnant in the direction of the constellation of Vela. This supernova remnant is only about 10,000 years old. It also contains a pulsar that has been observed to glitch.

No one has seen a pulsar in the process of glitching. Rather, the pulsar is observed at one time and then a little later, and the period is found to be slightly shorter. From such observations a few days apart, one can conclude that the glitches happen on a time that is shorter than a few days (possibly much shorter), but no more accurate statement can be made. The thing that is of particular interest is that after a glitch, the pulsar requires a considerable time, of order a month, to return to its original period and resume the same gradual lengthening of the period. That the time to return to normalcy is so long seems to strongly suggest that the inner portions of the neutron star are superfluid.

Glitches are thought to occur as a neutron star adjusts itself to the loss of rotational energy as it slows down. The understanding of how that adjustment occurs has evolved over the decades since glitches were discovered. An early model envisaged the spinning neutron star to form an equatorial "bulge" that was frozen in when the neutron star cooled and its outer layers solidified. As the neutron star spun more slowly, the bulge would settle by cracking and breaking. Conservation of angular momentum would cause the neutron star crust to rotate slightly more rapidly when the crust broke and settled into a smaller radius. This was thought to represent the formation of the glitch. The slow healing time was then thought to represent the long time necessary for the outer solid crust to bring the inner, zero viscosity, superfluid core into a common spin rate, after which the whole neutron star would begin to lose rotational energy and once again begin to spin ever more slowly. The reason to mention this picture is that it is a simple physical one that was reasonably easy to describe in lectures. I used it for decades, and it appears in other books. It is also wrong. More careful study showed that the mechanism of glitches is more interesting and subtle. The idea of crust cracking has survived in another context that will be described in Section 10.

The current model for glitches is based on considerations of exactly how the magnetic field that is such an obvious part of the external aspects of a pulsar threads the inner superfluid core. It turns out that a magnetic field cannot penetrate the superfluid, but only normal matter. For the magnetic field to thread the superfluid core, there must be "vortices" of normal matter that extend through the superfluid core, roughly parallel to the spin axis of the neutron star. The spinning vortices of normal matter are the repository of the angular momentum of the material in the inner core. The vortices of normal matter also provide the path for the magnetic field to pass from the north to the south pole within the neutron star. The vortices that allow normal matter and magnetic field to thread the superfluid are "pinned" to irregularities in the normal matter of the outer crust. In this picture, a glitch occurs because the vortices have a memory of the past when the outer crust was spinning faster. At intervals, some of the vortices unpin from the crust and coalesce, allowing the whole neutron star to adjust to its slower rotating, lower angular momentum state. Although the whole neutron star adjusts to the lower rotational state, this unpinning causes the outer crust to temporarily rotate more rapidly, giving rise to the glitch. As the neutron star attains its new equilibrium rotational state, the vortices again pin to the crust and slow it down so that the gradual slowing of the whole neutron star can continue. The bottom line is still that the glitch phenomenon cannot be explained without invoking a superfluid core.

5. Binary Pulsars – Tango por Dos

The accurate periods of pulsars make excellent clocks. If the clock were to move, the *frequency* of the pulses would be changed by the *Doppler shift*. The frequency of the radio emission would also be changed, but the radio radiation is continuum radiation, which, without spectral "lines," specific identifiable frequencies, gives no detectable Doppler shift. The pulses themselves are a marvelous substitute. With this clock, astronomers can look for periodic changes in the velocity of a pulsar that would indicate that the neutron star was in orbit. The evidence shows that to a high degree of accuracy the vast majority of pulsars are not in binary star orbits. Astronomers were very excited when in 1975 careful searches paid off, and a radio pulsar was discovered to be in a binary orbit. Since then, eleven more binary radio pulsars have been discovered. They are the exception that proves the rule: the vast majority of the known pulsars are single stars.

The discovery of the first binary pulsar led to a host of interesting results. The orbit was worked out from the Doppler shift of the pulsar period, and the prediction was made that any companion star of ordinary size would cause the eclipse of the neutron star once each orbit. No eclipse was seen. The lack of an eclipse implies that the companion star is itself a compact star, probably a white dwarf or neutron star.

Nature has been kind to put neutron stars in binary orbits. Study of the binary orbits allows the determination of the neutron star masses, a fundamental property that cannot be accurately measured by any present techniques for the multitude of single radio pulsars. The period of the orbit gives information about the masses of the stars, using Kepler's third law. The mass of the first binary pulsar is one of the few known neutron star masses. Both stars seem to have a mass of very nearly 1.4 solar masses. Other binary neutron stars have also had their masses weighed, and they also appear to have very nearly this mass. The coincidence of this number with the Chandrasekhar limit requires some comment. If a white dwarf attained the Chandrasekhar limit and collapsed to form a neutron star, the neutron star would be somewhat lower in mass. This is because some energy is inevitably ejected in the process of forming the neutron star, if only in the form of neutrinos. A great deal of energy must be ejected and the mass equivalent, in terms of $E = mc^2$, of the minimum energy loss is about 0.2 solar mass. To make a neutron star of 1.4 solar mass, the initially collapsing object would have to be 10 or 20 percent more massive, and hence somewhat greater than the Chandrasekhar mass. Just why neutron stars should form from cores of a precise mass that somewhat exceeds the Chandrasekhar mass is not clear.

The accurate orbital timing of the first binary pulsar showed that the orbit was decaying. The two stars are slowly spiraling together. Recall the final evolution of two white dwarfs from Chapter 5. They are imagined to spiral together as they give off gravitational radiation. In the binary pulsar system, the change in the orbit is precisely what would be predicted as the result of gravitational radiation. With one stroke, this observation confirms, indirectly but strongly, the predicted existence of gravitational radiation by Einstein's general theory and shows that gravitational radiation works in binary systems to draw stars together, just as the astrophysicists had predicted. Whatever the companion of this binary pulsar, white dwarf, or neutron star, gravitational radiation will eventually cause them to collide and merge. The discovery and analysis of the binary pulsar and the remarkable proof of gravitational radiation led to the award of the Nobel Prize to Joe Taylor and Russell Hulse, the radio astronomers at the University of Massachusetts (now both at Princeton) who made the discovery and analysis of the first binary pulsar. For this second Nobel Prize for work on neutron stars, the important contribution of the graduate student (Dr. Hulse) was recognized.

The binary pulsars, by being the exception to the rule, also lead us to ask why the strong majority of pulsars are not in binary systems. The binary pulsars provide a clue to the answer. One possibility is that neutron stars are commonly ejected from binary orbits by the explosion that creates them. Arguments based on conservation of energy and angular momentum show that if half the total mass of a binary system is ejected in an explosion, the system will be disrupted with the two stars flying off in opposite directions. In addition, pulsars are observed to sail through space at rather high velocities. There are a number of reasons to think that pulsars are given a "kick" by the process of violent gravitational collapse that creates them. Such kicks will also help to tear neutron stars away from any binary companion. Ejecting matter in the explosion and kicking the pulsars probably account for most of the single pulsars. The exceptions can also be understood at some level. For one thing, the star that blows up will frequently be the less massive star because it will have transferred mass to the companion. If the exploding star contains less than half the total mass of the two stars combined, then it cannot eject more than half the total mass, and the binary system will not be disrupted. The kicks to newly formed neutron stars may not be delivered in random directions, but inasmuch as they are, some of the kicks could help to keep the neutron star in orbit despite the loss of mass and gravity from the binary system by the supernova process itself.

The circumstantial evidence that Types Ib and Ic supernovae arise from massive stars that have lost their outer envelopes by mass transfer suggests that they create neutron stars in binary systems. Whether these neutron stars remain in the binary is not clear. There is a strong suspicion that, for systems in which the neutron star is still in a binary, the neutron star was born in some version of a Type Ib or Type Ic supernova explosion.

There may be another reason why the radio pulsars, in particular, are mostly single. An important feature of the first binary pulsar is that the companion star is known to be compact. No mass is being transferred in the system. As we will see in the next section, neutron stars are known to exist in binary systems for which the neutron star is not a radio pulsar. These systems are transferring mass. One reasonable hypothesis is that mass transfer prevents the emission of radio pulses by blocking the radio emission or by shorting out the sparking mechanism and preventing the radio radiation in the first place. With this picture, one would say that the binary pulsar is special not because the neutron star remained bound in a binary system, but because the companion star is unable to transfer mass and spoil the radio pulses. Those neutron stars that were always single stars or that were ejected from binary systems have no problem because they have no companion to interfere. Most neutron stars left in binary systems are not radio pulsars because they have the misfortune to be neighbors to a living star that insists on sharing some of its matter.

6. X-Rays from Neutron Stars – Hints of a Violent Universe

X-ray observations have been mentioned where appropriate throughout this book. The next subject owes its very existence to the advent of X-ray astronomy, however, and so a word of history is in order. In the last three decades, the science of X-ray astronomy has matured to become a major independent branch of astronomy. X-rays must be collected above the absorbing shield of the Earth's atmosphere. The first observations were made with brief rocket flights that only tantalized the scientists that launched them. There were glimpses of intense sources of high-energy X-rays.

The revolution in X-ray astronomy began with the launch of a small astronomical satellite dedicated to the detection of X-rays in 1972. The satellite was launched from a site in Kenya and was called *Uhuru*, the Swahili word for freedom. This first satellite could not locate the source of any X-ray emission very accurately, and, although better than rockets, it was not tremendously sensitive. *Uhuru* was on station for a long time compared to a rocket at perigee, however, and it could look for X-rays for orbit after orbit. The result was stupendous. The whole sky was alight with X-rays. It was like Galileo's invention of the telescope: to look with a new tool and to find that previously unknown or inconspicuous objects glared forth when examined properly. X-rays were seen from stars, from galaxies, from every direction! Above the protective layer of the atmosphere, the Universe was a far more violent place than astronomers had suspected.

Many X-ray satellites have been flown in the last 30 years. Several have been launched by the United States, others by European countries. Japan has had a very successful series of satellites and nearly took over the field when the U.S. support for X-ray astronomy lagged in the 1980s. Russia has also had a number of successful experiments. A major step of this first burst of activity in a new field was the launching by NASA of a large satellite in 1978 bearing the name *Einstein*, because it was the centennial year of his birth. This satellite contained a device that could focus X-rays like a proper telescope. It could measure details in an X-ray picture with an accuracy of 1 arc second, equivalent to that of ground-based optical telescopes. In 6 years, the science of X-ray astronomy made an advance in sensitivity and detail equivalent to the leap from Galileo's first telescope to the giant modern reflectors. The new *Chandra Observatory* mentioned in Chapter 1 is the latest step in this progression, and there are more and better projects under construction and on the drawing boards.

One of the subjects to benefit most greatly from the new science of X-ray astronomy was the study of neutron stars. This is because the great gravity of these objects causes tremendous heating of any matter that falls upon them. The matter becomes so hot that the maximum intensity of radiation comes in the X-ray portion of the spectrum. Under proper circumstances, neutron stars are just natural X-ray emitters.

Some of the first X-ray sources examined with *Uhuru* showed a peculiar behavior. The intensity of the X-rays was not constant, but faded away at regular intervals, typically every few days. Most of the scientists who worked on the early X-ray experiments building the detectors were physicists, not astronomers. The erratic behavior in the signal puzzled them. Astronomers – at least many amateurs who delight in such things, if not the professionals who specialized elsewhere – would have immediately identified the cause. The problem was that the X-rays were being eclipsed. The X-ray source was in a binary star orbit and was simply disappearing behind the other normal star periodically. This companion star was the source of matter that fell onto the neutron star and produced the X-rays.

This understanding led to a rapid series of identifications of orbiting neutron stars. A major new branch of astronomy was born almost overnight as the new sources were identified and characterized, and theorists rushed to understand their properties. The X-ray observations provided an exciting new way to probe the nature of mass transfer, accretion disks, and the structure and behavior of the neutron stars themselves. Although the existence of accretion disks had been demonstrated in the cataclysmic variables, it was the exciting new realm of neutron star X-ray sources that resulted in the sudden growth of interest and developments in the understanding of accretion disks.

Over the next few years after the launch of *Uhuru*, X-ray astronomers realized that there were two basic classes of binary neutron star X-ray sources (and a handful of oddballs that resist categorization). The first class consists of a neutron star in orbit about a normal, fairly low-mass star. The other class consists of neutron stars in orbit around high-mass normal stars. In this case, the normally evolving star typically has a mass in excess of 10 solar masses.

The classic example of the first type is the first X-ray source discovered by *Uhuru* in the direction of the constellation Hercules, the system named Hercules X-1. Detailed studies over decades have shown that Her X-1 is a nearly textbook example of mass transfer to a neutron star in a binary system, as shown schematically in Figure 8.2. A star of about 2 solar masses, slightly evolved on the main sequence, is filling its Roche lobe and transferring mass. The mass settles into an accretion disk. As friction operates in the disk, the matter spirals down toward the neutron star and gets heated. In the inner portions of the accretion disk, the orbital velocities are very high, so the frictional heating is strong, and the material in the disk itself emits X-rays. As the spiraling matter gets near the neutron star, the magnetic field of the neutron star channels the matter toward the magnetic poles. As the material finally lands on the surface of the neutron star, the impact causes more heating and further X-rays.

Although X-ray satellites are crucial to the discovery of X-ray sources, one should not forget that the astronomy advances most efficiently where standard earthbound optical techniques can be brought to bear in complementary studies. This is because, as a matter of practice, there is a tremendous amount of information available in the photons emitted in the optical band. This is, after all, where most stars emit the majority of their radiation. Most of our practical knowledge of the Universe is obtained in the optical, so X-ray (or radio, infrared, or ultraviolet, or gamma ray) information must be integrated into the realm of classical optical astronomy to come to full fruition.

As an example, studies of Her X-1 would be woefully incomplete without the optical studies of the companion star. It is the optical studies that tell us the type of star, its evolutionary state, and the fact that it is filling its Roche lobe. Coupled optical and X-ray studies were used to completely characterize the orbits of the two stars and to obtain a direct measure of their masses using Kepler's law. The mass of the neutron star comes out to be very nearly 1 solar mass. This mass seems to be significantly less than the 1.4 solar masses that has been measured so precisely for several of the binary pulsars, as mentioned in Section 5. There is no understanding of why this should be so. It is presumably an accident of birth of an especially low-mass progenitor core, but it might have involved an especially large ejection of the mass from the collapsing core. In this game, even “typical” objects are not so typical.

The observations of Her X-1 suggest that a star of initial mass between 10 and 15 solar masses evolved and shed its envelope. The bare core probably evolved on its own for a while and then collapsed. Like cataclysmic variables, there is a strong hint in Her X-1 that the original evolution was not just a simple case of one star losing mass to the other. For one thing, the two stars are too close together now for the first star to have developed a dense core and red-giant envelope. Also, the relatively low mass of the companion star suggests that it did not accept all the mass that the first star lost. Her X-1 is probably another example of common-envelope evolution in which the 2-solar-mass star was engulfed in the envelope of the more massive star. Much of the first star's envelope was presumably lost out of the system, and the core of the massive star and the smaller-mass companion spiraled together. Perhaps the smaller star filled its Roche lobe while still enshrouded in the envelope of the other. Whether any of this helps to explain the relatively low mass of the neutron star is not clear.

The other kind of binary X-ray source systems, those with high-mass normal companions, is typified by the third X-ray source *Uhuru* discovered in the direction of the constellation of Centaurus, Centaurus X-3. The basic difference between Her X-1 and Cen X-3 is that the mass-losing star in the latter is fairly massive, about 20 solar masses. This turns out to make an important modification to the mass transfer process, if not the ultimate outcome. When Cen X-3 was first discovered and the companion optical star identified, attempts were made to work out the orbits. According to the standard picture, the assumption was made that the companion filled its Roche lobe in order to transfer mass to the neutron star. The answers that emerged did not make sense. The mass of the neutron star was derived to be so low, about 0.1 solar mass, that the gravity should be so weak that any neutron star should expand to be a white dwarf instead.

The problem was that the companion star does not fill its Roche lobe! Rather, such a massive star blows an appreciable stellar wind. It loses mass through this wind whether it has a companion star or not. In this case, however, there is a neutron star the gravity of which reaches out and ensnares some of the passing wind. The matter from the wind then settles into an accretion disk. With this picture, things make more sense. The orbital information from Cen X-3 is not as accurate as that from Her X-1, never mind the binary pulsars. The best estimate for the mass of the neutron star comes out to be a little more than a solar mass, but a mass of 1.5 solar masses cannot be excluded. This is a reasonable result.

The disproportionate mass between the neutron star and the massive normal companion in Cen X-3 has one interesting consequence. The neutron star raises tides on the surface of the companion, just as the Moon does on the Earth. Energy is expended in dragging those tides around, and the energy comes out of the orbit, causing the neutron star to spiral toward the other star. If the companion is not too massive, the tidal drag causes it to spin faster until the companion rotates at exactly the speed that the neutron star orbits. Then the tide just sits in one place on the surface of the star, and there is no drag. For a massive companion, however, there is too much inertia. The central star and the tides always lag behind the orbital motion, dragging the neutron star down. There is no limit to this process, and eventually the neutron star should collide with and disappear into the companion star. The neutron star could spiral to the center, swallow matter from the star, collapse to make a black hole, and then eat the whole star! This may be the fate in store for Cen X-3.

Her X-1 and Cen X-3 share another very important feature. The X-rays they emit come in pulses, 1.2 seconds apart for Her X-1 and 4.8 seconds for Cen X-3. The behavior is very reminiscent of the pulses from radio pulsars, but the energy is coming in the X-ray portion of the spectrum. In addition, for extended periods of time the pulses get steadily more rapid, whereas, except for glitches, the radio pulses slow down.

Despite the exotic nature of the radiation, the X-ray pulses are easier to explain than the radio pulses. Much of the explanation borrows heavily from the knowledge gained by studying radio pulsars. The neutron stars are presumed to be magnetized and rotating. The crucial difference is that, whereas a pulsar must generate radio radiation by its own devices, the X-rays are caused by an external agent, the dumping of mass upon the neutron star.

With the presence of the magnetic field, the matter arrives at the neutron star in a special way that promotes pulses. The matter spirals down in the accretion disk until it encounters the outer reaches of the magnetic field. At that point, the matter finds that it cannot continue in orbit because it cannot move across the lines of magnetic force. Rather, the matter falls along the lines of force, as shown in Figure 8.2. These lead naturally to the north and south magnetic poles of the neutron star. The matter is channeled so that it falls selectively on the magnetic poles, not at random on the surface of the neutron star. The intense X-radiation then comes from the magnetic poles, as if there were two bright spots on an otherwise dark surface. If the magnetic axis is misaligned with the axis of rotation, then, as the neutron star spins around, first one then the other bright spot points at the Earth, just like a lighthouse. The observer detects a pulse of X-rays as the pole is swept into view by the rotation. With mass transfer, one can understand fairly easily why the radiation comes from the poles and hence why there are pulses.

The influence of mass transfer also explains why the pulses tend to speed up rather than slow down. There are two competing effects. The loss of energy in the radiation tries to slow the neutron star down. The matter arriving from the accretion disk, however, carries with it the angular momentum of its orbit. As the matter lands on the neutron star, the spin is transferred to the neutron star. This turns out to be the dominant effect in many circumstances, and the neutron star rotates faster and faster until the mass transfer stops or the neutron star is rotating as fast as the accreting matter where it begins to interact with the magnetic field. If the neutron star tries to rotate too fast, its magnetic field acts like a paddle to splash matter out of the accretion disk, which slows the neutron star down. Both Her X-1 and Cen X-3 have gone through episodes lasting a couple of years where they have stopped speeding up (Cen X-3) or have even tended to spin more slowly (Her X-1). This is presumably because they have ejected matter or the rate of mass transfer has declined so the accretion disk has retreated, allowing the neutron star rotation to slow. Even though the spin-up by accretion makes good sense, the slow-down process must be rather prevalent because many X-ray pulsars have rather long periods, some as long as 800 seconds.

7. X-Ray Flares – A Story Retold

Recall from Chapter 5 that there were two basic classes of flaring binary white dwarf systems: the dwarf novae where the accretion disk is the source of the activity and classical novae caused by thermonuclear explosions on the surface of the white dwarf. Suppose the white dwarf were replaced by a neutron star. Similar phenomena will occur.

X-ray astronomers see several accreting neutron stars in the Galaxy that are labeled as *X-ray transients*. In this context, the general word “transient” refers to a particular phenomenology, implying a particular physical cause. Every few years, these X-ray transients emit a flare of X-rays that lasts for about a month or so. At least two of these systems are well studied and are known to be in binary systems. There is a strong suspicion that the process causing this outburst is similar to that in dwarf novae, an instability in the flow in the accretion disk. The accretion disk instability described in Chapter 4 does not depend sensitively on the nature of the object around which the disk circles. If matter flows into the disk from a companion star at an appropriate rate, the disk will go into the storing and flushing mode that characterizes the dwarf novae. If the object receiving the mass is a neutron star, however, then in the flushing phase, matter from the disk is spiraling down onto a neutron star. The matter gets intensely hot and emits X-rays. The time scales are somewhat longer in the X-ray transients than in dwarf novae, and there are no quantitative models, but the disk instability is a plausible picture for the origin of the X-ray transients.

There is also a neutron star analog of classical novae. In 1978, a fascinating new class of X-ray sources was discovered. Russian scientists first noticed the phenomena. Some X-ray sources show an occasional brief, strong burst. The power rises in about a second and then decays over the course of the next minute or so. The bursts recur every few hours more or less randomly. After the Russians reported these bursts, a search of old *Uhuru* data also showed the effect. The American astronomers just had not noticed it at first in the welter of data with which they had to deal.

The display in the X-ray bursts is not like the rather demure pulses from Her X-1 and Cen X-3 or like the occasional flares of the X-ray transients. The bursts are very energetic compared to the pulses of Her X-1 or Cen X-3. They are comparable in power to the X-ray transients but much shorter in duration. They call for a completely different physical explanation.

Of the more than 100 X-ray sources in the Galaxy with low-mass companions, about 40 are *X-ray bursters*. None of the binary X-ray sources with high-mass companions display this behavior, and neither do the few low-mass systems that display X-ray pulses like Her X-1. Like the general population with low-mass companions, the X-ray bursters tend to cluster toward the center of the Galaxy, as do the oldest stars in the Galaxy. At least nine of the X-ray bursters are seen to be in globular clusters that are also old assemblages of stars. Most X-ray bursters show no evidence for binary motion, but evidence has been reported for orbital motion in at least one X-ray burst source. The guess is that all these systems are in binary systems, but nature conspires to hide the fact. If the systems are seen edge-on, it is most easy to determine the Doppler motion due to their orbit, but in this case the neutron star and its X-rays can be obscured by the accretion disk. If the system is nearly face-on, the X-rays can be seen, but the orbit is difficult to determine because all the motion is almost at right angles to the observer. The Doppler shift only registers the component of motion directly toward or away from the observer. The X-ray bursters do not show any sign of X-ray pulses (an exception will be described later). The interpretation is that the neutron stars in these systems have very low magnetic fields, so matter is not focused on the magnetic poles, and there is no X-ray "lighthouse" effect.

The theory for the burst sources is based on thermonuclear explosions on the surface of the neutron stars. Calculations have shown that as hydrogen accretes onto the surface of a neutron star, it is heated and burns in a regulated fashion. Under proper circumstances, the resulting helium, however, piles up in a layer supported by the quantum pressure. As we have seen in several instances, this condition leads to unstable burning when the helium finally gets hot and dense enough to ignite. The X-ray bursts are thus thermonuclear explosions on the surfaces of the neutron stars. There is therefore a direct parallel for this explanation for the X-ray bursts and the explanation for the outbursts in the classical novae, the basic differences being in the nature of the compact object doing the accreting. Because of the high gravity of neutron stars, relatively little, if any, matter is ejected from the neutron star in an X-ray burst. The high gravity also causes the very short time scale of the explosion on the surface of the neutron star as compared to the effects in a classical nova that can linger for a year or more.

The theory of these nuclear outbursts shows that they only occur if the rate of accretion of matter onto the neutron star is relatively sedate. This allows the layer of helium to build up supported by the quantum pressure. At high accretion rates, the helium stays hot, is supported by the thermal pressure, and burns in a regulated, nonflaring way. One of the implications of this theory is that if the neutron star is strongly magnetic, then even a sedate rate of accretion will be focused onto the magnetic poles, giving an effective high rate of accretion at those two spots. That will give the circumstances for hot magnetic poles and X-ray pulses, but it will mean that the rate of the accretion at the poles is high enough that the helium will ignite and burn in a regulated way. This is another reason to argue that neutron stars that show X-ray pulses have large magnetic fields and no X-ray bursts and neutron stars that show X-ray bursts have small magnetic fields and do not display X-ray pulses.

The Eddington limit discussed in Chapter 2 plays an interesting role in the neutron star accretion process associated with these X-ray burst sources. Recall that the Eddington limit is a limit to how bright an object can be without blowing away matter by the sheer pressure of the outflowing radiation. The Eddington limit depends on the gravity of the object, and so the limiting luminosity scales with the mass. For accreting neutron stars, there is a close coupling between the mass and the luminosity because the luminosity is caused by the infalling matter. This means that if the matter falls in at too high a rate, intense radiation will be generated. The infalling matter will be blown away rather than accreting. If too much of the infalling matter is blown away, however, then there is not enough radiation to blow the matter away, and the infall can take place. The result can be to balance things so that some matter is blown away and some accretes. The luminosity adjusts so that the Eddington limit is not violated. Many of the binary neutron star X-ray sources have luminosities somewhat below the Eddington limit, as if they had made their accommodation with the limiting luminosity. In the observed X-ray bursts, the luminosity rises until it bumps right into the ceiling of the expected Eddington limit for an object of the mass of a neutron star, about 1 solar mass.

At least one binary neutron star system, Centaurus X-4, displays both X-ray transient outbursts and X-ray bursts. As the X-ray flux from Centaurus X-4 declined from one month-long flare of the X-ray transient variety, it showed another brief flare of the X-ray burst variety before proceeding to decline. Presumably an accretion disk instability flushed matter down toward the neutron star creating the X-ray transient. As matter accumulated on the neutron star, it underwent a thermonuclear outburst. Then the disk went into its storage mode; there was no fresh mass added to the neutron star, so no repeated X-ray burst.

8. The Rapid Burster – None of the Above

One particular source, the *Rapid Burster*, displays behavior that falls in the “none of the above” category. This system, known to intimates as MXB 1730–335 (for MIT X-ray Burst), was discovered about 20 years ago. When active, it bursts about four thousand times a day. The Rapid Burster is located in a globular cluster. It also occasionally has the more prominent bursts associated with the thermonuclear ignition of helium. Like the other thermonuclear burst sources, the Rapid Burster shows no sign of X-ray pulses that would indicate the rotation of the underlying neutron star. The presumption is that the magnetic field of this neutron star is relatively weak so matter falls more uniformly on the surface and is not focused at the magnetic poles. The repetitive bursts that define the Rapid Burster are thought to be neither a thermonuclear burst on the surface of the neutron star nor the type of accretion disk heating instability similar to that of dwarf novae. The observations suggest that the matter rains down on the neutron star in blobs like a rapidly dripping faucet, rather than in a steady gush. There is no well-established theory for this behavior, but the suspicion is that it involves an instability of the matter on the inner edge of the accretion disk that may be due to a condition where the pressure of radiation becomes excessively large, larger than the pressure of the hot gas in the disk. For 20 years, the Rapid Burster was alone, but now it has some company.

In 1990, NASA launched another of its great observatories to complement the *Hubble Space Telescope*. This was the *Compton Gamma-Ray Observatory*. We will talk about it more in Chapter 11. In December 1995, this satellite discovered a system known as the *bursting pulsar*, or, more technically according to its discovery instrument and coordinates, GRO J1744–28. Follow-up work on it was done by another NASA satellite, the *Rossi X-ray Timing Explorer*. This relatively modest satellite was named after Bruno Rossi, an MIT pioneer of X-ray astronomy, and was designed to follow X-ray behavior with very accurate timing. Observations with *RXTE* of the bursting pulsar showed an incredible array of behavior that indicate that this system may be an important link between systems like the Rapid Burster and the other X-ray burst sources.

As its name implies, the bursting pulsar is an X-ray pulsar. From the frequency of the pulses, one can deduce that the neutron star rotates about twice a second. Its orbital motion has also been detected. The neutron star is in a 12-day orbit around a small red giant that has lost almost all of its hydrogen envelope and now has a mass of about one-quarter the mass of the Sun. From January through May of 1996, the bursting pulsar showed large bursts lasting about 10 seconds apiece every 2 hours or so. These bursts displayed characteristics of the staccato bursts of the Rapid Burster rather than the helium ignition flares of the X-ray bursters. The presumption is that the bursting pulsar has a stronger magnetic field than the Rapid Burster and hence can both generate “lighthouse” pulses of X-rays from the magnetic poles and can suppress nuclear flares by the focused, hot accretion at the magnetic poles. The fact that it still manages to show the instability of the inner disk means that the magnetic field is not so strong that it cuts out the inner part of the disk where that instability happens. The bursting pulsar is thus an interesting intermediate case that promises to teach us more about the conditions under which neutron stars evolve in binary systems. After May of 1996, the system got so dim that *RXTE* could no longer detect it, so for now, the bursting pulsar is keeping any further secrets it may have to reveal.

9. Millisecond Pulsars

In the last decade, a new variety of radio pulsars have been found that have generated great excitement because they link so many aspects of the formation and evolution of neutron stars. Theory predicts that neutron stars cannot rotate faster than about one thousand times per second without flinging themselves apart with the excessive centrifugal force. That limiting rotation rate corresponds to a rotational period of 0.001 second, or 1 millisecond. Thus one expects that the fastest pulses that could be discovered from a pulsar would be about 1 millisecond, and that a 1 millisecond pulsar would be on the verge of tearing itself apart. Realistically, one would expect that pulsars would rotate a little slower than this fastest possible limit and, hence, to have pulses of a few milliseconds. By this standard, the pulse period of the Crab nebula pulsar is dawdling along at a mere 33 milliseconds.

Special search techniques were developed to search for pulsars near this period limit, and they have been successful. Over two dozen *millisecond pulsars* have been found. In contrast to their longer period kin, about half of the millisecond pulsars are in binaries. The most rapidly rotating has a remarkably well-defined period, 0.00155780644885 seconds, or about 1.6 milliseconds. This neutron star is whipping around 642 times per second.

The next step is to account for the origin of the millisecond pulsars. Pulsars must be magnetic neutron stars. The Crab pulsar rotates 30 times per second; normal pulsars, about once per second. This is because the Crab pulsar is only 1,000 years old. When it is several million years old, the Crab pulsar will have slowed down, and it will presumably also have a period of about 1 second. This suggests that millisecond pulsars might be very young, newly born neutron stars. More thought, and appropriate observation, shows just the opposite is the case. With a normal-strength magnetic field, a pulsar with a period of 1 millisecond would be losing energy so fast that it could not maintain its rapid rotation. By this argument, the millisecond pulsars should be slowing down very rapidly, but they are observed to be slowing scarcely at all. The millisecond pulsars must therefore have a smaller magnetic field than normal so that they lose little rotational energy into radiation. This in turn suggests that they are old so that there has been time for their magnetic fields to decay away or otherwise disappear. If they are old, however, why have they not lost more of their rotational energy when they were younger with a more robust magnetic field?

The proposed resolution to this query is that the neutron stars were born in binary systems and that transfer of mass and associated angular momentum from a companion kept the neutron star spinning fast even as the field decayed. Thus all millisecond pulsars should be in binary systems, but a significant fraction of them are not. This is another dilemma. If there were a binary companion, where did it go?

One possible answer to this further dilemma was suggested by the discovery of a particular millisecond pulsar in a binary system. This pulsar orbited a companion, a more or less normal star. It appeared as if the pulsar were killing the normal star because the star was losing mass at a high rate. The rapidly rotating neutron star produces a great flux of high-energy radiation, X-rays and gamma rays. It was first thought that this intense radiation was literally blasting away the companion star. Some astronomers termed this system the Black Widow star because the neutron star was perceived to be killing its mate. Subsequent observations showed that the star was probably losing mass on its own. In any case, the implication is that the companion will soon be gone, leaving a millisecond pulsar to spin alone in space. Roughly half of the millisecond pulsars are in binary systems with a companion star to transfer mass and keep them spun up. Presumably the other half of the observed millisecond pulsars have already dispensed with their companions in one way or another.

Another interesting millisecond pulsar revealed that it had objects of planetary mass orbiting it. These objects were discovered only by the exquisite timing that is possible with these pulsars. Tiny rhythmic oscillations in the pulse period revealed that the pulsar was being slightly tugged around in space by several small objects of mass about that of Jupiter. Whether these are true planets, left over from some ill-fated solar system that orbited the star before it exploded, or whether the “planets” are themselves left over lumps of blasted star-stuff is not clear.

To put the millisecond pulsars in perspective, we need to take a step back in the evolutionary story. What sort of system gave rise to the original system of a neutron star orbiting an ordinary star? The explosion of a supernova in a binary system ejects a great deal of mass and hence decreases the gravity that holds a binary system together. That is why we think most ordinary pulsars are alone in space. They have not murdered their companions, but they may have unbound and ejected them from orbit. To prevent this, we need a fairly gentle way to make a neutron star. After the neutron star is born, it must have a weak magnetic field or lose an originally strong magnetic field and then be spun up by accretion to become a millisecond pulsar.

If this is the evolution of the neutron stars that become millisecond pulsars, then such systems should pass through a phase in which the companion adds mass to the neutron star to spin it up. The result should be the production of X-rays. The natural conclusion is that the systems we see now as X-ray sources with neutron stars orbiting low-mass companion stars will evolve to become the millisecond pulsars. The problem is that if you work out the rate at which X-ray systems with neutron stars and low-mass companions are born and the rate at which millisecond pulsars are born, they disagree substantially. There do not seem to be enough low-mass X-ray systems to account for the number of millisecond pulsars. Either there is another way to make millisecond pulsars, or there is something we do not understand about the evolution of the stellar systems in the X-ray phase. If that phase lasted a shorter time than we think, there would have to be a higher production rate to account for the number we see at this epoch in galactic history. That would help close the gap.

Another mechanism that might avoid the phase of being an ordinary X-ray source during the spin-up phase has been suggested to produce millisecond pulsars. That mechanism involves the accretion of matter onto the O/Ne/Mg core of a star of original mass of about 10 solar masses. When such a core reaches its maximum mass, it will undergo electron capture and collapse to form a neutron star, but essentially all the core will collapse to make the neutron star, and very little is expected to be ejected (Chapter 6). This gives the maximum probability of maintaining a companion in binary orbit. This general process is called *accretion induced collapse*, to distinguish it from core collapse brought on by the normal process of core collapse of a single evolving star as fuel is burned to heavier elements. This process is plausible in general, but it does not necessarily predict that the resulting neutron star will be rapidly spinning with a low magnetic field, the conditions required to be a millisecond pulsar.

The low magnetic fields required to explain the millisecond pulsars have raised a different conundrum. All radio pulsars are observed to fall on the short period side of a limiting value of the period that depends on the strength of the magnetic field. The implication is that as pulsars age and rotate slower and slower, their magnetic fields decay away so that for very old slow pulsars the combination of rotation and magnetic field is no longer able to generate the thunderstorms at the magnetic poles that are required to make a radio pulsar. In a plot of magnetic field versus spin period, this limiting period is known as the “death line.” Taking a somewhat more pragmatic approach, Mal Ruderman of Columbia University argues that the cutoff may be different for different magnetic field configurations and hence the boundary may be a “death valley.” In any case, the notion persisted for two decades that the magnetic field of pulsars decays away with a time scale of perhaps 100 million years. Continued consideration of the numbers of pulsars with different field strengths and spin periods and the existence of the millisecond pulsars with very low magnetic fields has inspired reconsideration of this issue. There are suggestions that the field may not decay or that it is the accretion process itself that kills the field in the case of the millisecond pulsars. The origin and evolution of neutron star magnetic fields is still a subject of active investigation.

10. Soft Gamma-Ray Repeaters – Reach Out and Touch Someone

Although the Sun occasionally belches a flare of particles that reach the Earth and affect radio communications, we are used to the stars being quietly remote in their isolated magnificence against the backdrop of dark space. Imagine our surprise, therefore, when one of them reached out and touched us in August of 1998! As the Earth sails around the Sun and follows the Sun around the Galaxy for billions of years, it is not isolated from the violent Universe around us.

A class of bursting events called *soft gamma-ray repeaters* has been defined over the last 20 years. At first these events were confused and intermingled with the events known as *gamma-ray bursts*, the story of which we will learn in Chapter 11. The difference between "hard," high-energy X-rays and "soft," low-energy gamma rays is a matter of operational definition, and the dividing line is somewhat arbitrary. As the names imply, however, soft gamma-ray repeaters and gamma-ray bursts radiate most of their energy in the gamma-ray range. The soft gamma-ray repeaters emit somewhat less energetic photons than the gamma-ray bursts, a difference an expert can love. As we shall see in Chapter 11, no gamma-ray burst has ever been known to repeat. As data accumulated, however, it became clear that the sources that gave out the softer gamma-rays could and did repeat their outbursts, if at irregular intervals. The question was, what were they? Gamma rays of any sort require high energies, and that suggests high gravity, so one might think about white dwarfs, neutron stars, or black holes. Round up the usual suspects! An important clue was that all the soft gamma-ray repeaters turned out to be in supernova remnants.

The currently most widely accepted theory for the soft gamma-ray repeaters was developed by Rob Duncan at the University of Texas and Chris Thompson now at the University of North Carolina. They were originally seeking an explanation for gamma-ray bursts, not soft gamma-ray repeaters. Their investigations led them to consider neutron stars with very strong magnetic fields. They developed a theory that, under certain circumstances involving, among other things, very rapid rotation, neutron stars could develop immensely strong magnetic fields. Whereas millisecond pulsars have magnetic fields about ten thousand times less strong than "normal" pulsars, Duncan and Thompson argued for magnetic fields thousands of times stronger than "normal." The force of such magnetic fields could rival the gravity of the neutron star – strong indeed. Duncan and Thompson needed a name to distinguish their intellectual baby from the "normal" pulsars and millisecond pulsars, so they coined the phrase *magnetar* for a neutron star where the magnetic field rivaled gravity and pressure.

As they investigated the properties of magnetars, Duncan and Thompson realized that they should have a special activity. When they are first born, the magnetars would assume an equilibrium balancing the magnetic fields, pressure, gravity, and the centrifugal force of their rapid rotation. The latter would cause the neutron star to bulge along the equator, and that bulge would tend to be frozen into place in the outer rocky crust of the neutron star. As the neutron star lost energy and slowed, the bulge would be too big for the slower rotation, and it would eventually crack and settle. This picture is very similar to the original explanation for glitches in pulsar rotation rates, which has now been supplanted, as mentioned in Section 4. In the context of the magnetar theory, however, Duncan and Thompson realized that such a crust cracking would send powerful waves into the magnetic field that looped above the neutron star surface. The magnetic field would have to readjust to the new structure of the neutron star, and the magnetic field would convert some energy into hot plasma. That hot plasma would radiate the gamma-ray energy for the time scales observed in soft gamma-ray repeaters. Duncan and Thompson proposed that soft gamma-ray repeaters were, in fact, magnetars, a variety of super magnetized neutron star not previously recognized. They also recognized that after the first, major, crust-cracking star quake, there could be more localized shifts in the crust as it adjusted to the rearranged magnetic field. This would give a smaller, dimmer source of soft gamma rays, but if the spot were carried around by the rotation of the neutron star, then one might see a “lighthouse” effect so that the gamma rays would be seen to “pulse” at the rotation rate of the neutron star.

This suggestion that soft gamma-ray repeaters were magnetars attracted some positive, some negative, and some bewildered reactions. To make progress, observational confirmation was needed, and that came in 1998 in a rapid succession of events. Careful observations with the *Rossi X-ray Timing Explorer* revealed the rotation rate and rate of slowing down of one of the soft gamma-ray repeaters. The observations were consistent with a neutron star with a magnetic field one thousand times stronger than “normal.”

In August of 1998, Nature made sure we understood this lesson. One of the soft gamma-ray repeaters went off with a burst that was so strong that it affected the Earth! The gamma rays from this soft gamma-ray repeater affected the ionization of the upper atmosphere and interfered with radio communications worldwide. A wonderful contribution to the Op/Ed page of the *New York Times* described the awe-inspiring, incredibly intense, and widespread aurora witnessed by a bunch of guys on a fishing expedition above the Arctic Circle. This was the first known event when a star beyond the Solar System physically affected the Earth. There was no harm done, but this cannot have been the first time such a thing happened, and it will not be the last.

The event just described also brought evidence for a pulsar with a superstrong magnetic field. It had the immensely strong burst that tickled the Earth's ionosphere, but then a series of pulses just as Duncan and Thompson had predicted for the subsequent hot spots that should occur as the crust shifted in spots. In hindsight, just this behavior had been seen in the first soft gamma-ray repeater observed in 1979 in the Large Magellanic Cloud. At the time, that outburst was strange and controversial. That misery is now comforted by the company of the nearly twin outburst of the nearby source that produced the August 1998 burst. One must be careful and continue to seek evidence, but the magnetar theory is clearly the leading contender to account for the soft gamma-ray repeaters.

There is a handful of other objects that also seem to fit nicely into this scheme. These have been known as the *anomalous X-ray pulsars*. Like the soft gamma-ray repeaters, the anomalous X-ray pulsars are all found in supernova remnants. They show no evidence for binary companions. They all have rather long periods that fall in a restricted range of 6–11 seconds, very similar to the soft gamma-ray repeaters. They all seem to be spinning down, the spin periods getting longer and longer as if the spinning source were simply losing energy. From the spin period and rate of decrease of spin, an indirect estimate can be made of the strength of the magnetic field and the result is a value comparable to magnetars: 100 to 1,000 times stronger than normal radio pulsars.

A scheme that makes sense is that one neutron star in ten is born with an especially high rotation that allows the newly born neutron star to generate the high magnetic field. For the next 1,000 years, that magnetar undergoes crust cracks and rearrangement and is active as a soft gamma-ray repeater. After that time, the neutron star rotates sufficiently slowly that it cannot generate strong gamma-ray outbursts, but for the next 40,000 years it can radiate enough to be seen as an anomalous X-ray pulsar. After that time, it will be cooler and slower and will be a “dead” magnetar. The nature of the supernova that gives rise to magnetars and the nature of dead magnetars are not clear. How often do we end topics on that note? Such a big Universe, so little time. . . .

11. Geminga

Yet another chapter in the neutron star story is told in the saga of the source known as *Geminga*. This source was first detected by one of the early satellites with gamma-ray instruments in 1973. Two decades were required to figure out what it was. The name was given to it by an Italian X-ray astronomer, Giovanni Bignami. The name is nominally related to the fact that it is a gamma-ray source in the direction of the constellation of Gemini. More amusingly, it is an Italian double entendre related to the fact that the source could not be detected in the radio, one of the on-going mysteries. In the dialect of Milan spoken by Bignami, *ghe'è minga* means it's not there.

Vision in gamma rays is blurry and there were lots of spots of light in the direction of Geminga. A long time was required to pin down the source. In the optical, stars, asteroids, and plate defects had to be ruled out. The *Einstein* satellite revealed an X-ray source that helped to narrow down the optical search. One thing became clear. Whatever the object was, it was damn dim in the optical. Suspicion that Geminga was a neutron star grew. In the late 1980s, a dim optical source was isolated. It turned out to be the real thing.

A major breakthrough came finally in the 1990s with observations from the *Compton Gamma-Ray Observatory* and the *German Röntgen Satellite* or *ROSAT*, named for the discoverer of X-rays. Observations with these instruments showed that Geminga revealed both gamma-ray and X-ray pulses due to rotation with a period of 0.237 seconds. Geminga was a neutron star. Like the Crab nebula pulsar it emitted gamma rays, but unlike the Crab pulsar and so many others, it did not emit radio radiation. Various arguments suggested that it was very close to the Earth. That meant that, even though the gamma rays were detected, they were intrinsically feeble. That was why similar sources were not common. They would just be too hard to detect at greater distances. The small distance also explained why Geminga could be seen in the optical at all. Neutron stars have such a small radiating surface that one would have to be very close to be observed.

The close distance had another significant implication. There was a chance to detect the *proper motion* of the source, the motion across the sky due to its motion through space, and even the *parallax*, the apparent motion due to the Earth's orbit around the Sun. The former gives a hint of where Geminga arose; the latter, how far away it is. The parallax was measured in 1994 with the *Hubble Space Telescope*, and Geminga is only 160 parsecs, about 500 light years away – right in our back yard! The proper motion was extrapolated backward, and Geminga's origin was traced to near a star in the Orion nebula. There is an expanding cloud of gas around a star there that might be the supernova that created Geminga. The time for Geminga to get from Orion to where it is now is about 350,000 years, which is consistent with the age measured from the rate of slowing of the spin and with the estimated age of the supernova remnant. There are other possible interpretations, but the strong implication is that Geminga arose in a supernova explosion rather nearby about 350,000 years ago. Early hominids were leaving the veldt then and beginning to explore the planet – not so long ago.

The interpretation of Geminga is that it is a neutron star with a rather normal magnetic field. In 350,000 years, it has spun down so that it can barely generate gamma rays by particle creation and acceleration near the magnetic poles. Its surface is still hot and glows in the optical, if dimly. The most likely reason why radio is not observed is that the radio is created, but that it is radiated away from the Earth by an accident of orientation. Overall, Geminga is very special because of its nearness to Earth, but it may represent a normal phase in the aging and evolution of normal neutron stars. In looking to the past of Geminga, we may also be looking to the future when Betelgeuse erupts at about the same distance, the story foretold in the sidebar in Chapter 6.