Physics – Uspekhi **37** (7) 697 – 704 (1994)

PHYSICS OF OUR DAYS

PACS numbers: 97.60.Gb; 04.30.-w; 01.65.g

The binary pulsar, gravitational waves, and the Nobel Prize

C M Will

It is unlikely that Joe Taylor and Russell Hulse will ever forget the summer of 1974. It started uneventfully enough. Taylor, a young professor at the University of Massachusetts at Amherst, had arranged for his graduate student Hulse to spend the summer at the Arecibo radio telescope in Puerto Rico looking for pulsars. They had put together a sophisticated observational technique that would allow them to scan a large portion of the sky, using the radio telescope in such a way that it would be especially sensitive to signals from pulsars. At that time, over 100 pulsars were known, so their main goal was to add new ones to that list, in the hope that — by sheer weight of number — they could learn more about this class of astronomical objects. But apart from the possible payoff at the end of the observations, the bulk of the summer would be spent in rather routine, repetitive observing runs and compilation of data, that, as in many such astronomical search programs, would border on tedium.

But on July 2, serendipity struck.

On that day, almost by accident, Hulse discovered something that would catapult them into the astronomical headlines, excite the astrophysics and relativity communities, and ultimately yield the first confirmation of one of the most interesting and important predictions of general relativity. In the end, it would send them to Stockholm and the 1993 Nobel Prize. At least as far as relativists are concerned, their discovery ranks almost up there with the discovery of pulsars themselves.

That discovery was equally serendipitous. In late 1967, radio astronomers Jocelyn Bell and Antony Hewish at Cambridge University were attempting to study the phenomenon of scintillation of radio sources, a rapid variation or 'twinkling' of the radio signal from these sources that is caused by clouds of electrons in the solar wind out in interplanetary space. These variations are typically random in nature and are weaker at night when the telescope is directed away from the Sun, but in the middle of the night of November 28, 1967, Bell, who at the time was a graduate student, recorded a sequence of unusually strong, surprisingly regular pulses in the signal. After a month of further observation, she and Hewish established that the source was outside the solar system, and that the signal was a rapid set of pulses, with a period of 1.3373011 seconds. As a standard of time measurement, these pulses were as good as any atomic clock that existed at the time. It was so unexpected to have a naturally occurring astrophysical source with such a regular period, that, for a while, they entertained the thought that the signals were a beacon from an extraterrestrial civilization. They even denoted their source LGM, for Little Green Men. Soon the Cambridge astronomers discovered three more of these sources, with periods ranging from one quarter to one and a quarter seconds, and other observatories followed with their own discoveries. The Little Green Men theory was quickly dropped, and the objects were renamed 'pulsars' because of the pulsed radio emission.

This discovery had a tremendous impact on the world of astronomy. The discovery paper for the first pulsar was published on February 24, 1968 in the journal *Nature*, and in the remaining 10 months of that year, over 100 papers were published reporting either observations of pulsars or theoretical interpretations of the pulsar phenomenon. In 1974, Hewish was rewarded for the discovery with the Nobel Prize in physics, along with Sir Martin Ryle, one of the pioneers of the British radio astronomy program. In some circles, controversy still lingers over the decision of the Swedish Academy not to include Ms Bell in the award.

Within a few years of the discovery, there was general agreement among theorists and observers about the overall nature of pulsars, although many of the details are still not completely ironed out. Pulsars are simply cosmic lighthouses: rotating beacons of radiowaves (and in some cases of optical light, X-rays, and gamma rays) whose signal intersects our line of sight once every rotation period. The underlying object that is doing the rotating is a neutron star — a highly condensed body, typically of about the same mass as the Sun, but compressed into a sphere of around 20 km in diameter, 500 times smaller than a white dwarf of a comparable mass. Its density is therefore about 5×10^{14} g cm^{-3}, comparable to the density inside the atomic nucleus, and its composition is primarily neutrons, with a contamination of protons and an equal number of electrons. Because the neutron star is so dense, it behaves as the ultimate flywheel, its rotation rate kept constant by the inability of frictional forces to overcome its enormous inertia. Actually, there are some residual frictional forces between the neutron star and the surrounding medium that do tend to slow it down, but an example of how small this effect can be is given by the original pulsar: its period of 1.3373 ... seconds is observed to increase by only 42 nanoseconds per year. Of the 100 or so pulsars known by 1974, every one obeyed the general rule that it emits radio pulses of short period (between

C M Will McDonnell Center for the Space Sciences, Department of Physics, Washington University, St. Louis, Missouri, USA
Tel. (314) 935-6244
E-mail: cmw@wuphys.wustl.edu

Uspekhi Fizicheskikh Nauk **164** (7) 765 – 773 (1994)
Submitted in English by the author

fractions of a second and a few seconds), and with a period that is extremely stable, except for a very, very slow increase. We will see that this rule almost proved to be the downfall of Hulse and Taylor.

Other aspects of pulsars were not (and are still not) so well understood, however, and one of these is the actual mechanism for the 'lighthouse beacon', if indeed that is how the radio pulses are produced. In the conventional model, a pulsar is thought to have one important feature in common with the Earth: its magnetic north and south poles do not point in the same direction as its rotation axis. There is one key difference, however. The magnetic field of a pulsar is 10^{12} times as strong as that of the Earth. Such enormous magnetic fields produce forces that can strip electrons and ions from the surface of the neutron star and accelerate them to nearly the speed of light. This causes the particles to radiate copiously in radio waves and in other parts of the electromagnetic spectrum. And since the magnetic field is strongest at the poles, the resulting radiation is beamed outward along the north and south magnetic poles. Because these poles are not aligned with the rotation axis the two beams sweep the sky, and if one of them hits us we call it a pulsar. The details of this mechanism are extremely difficult to work out, partly because we have absolutely no laboratory experience with magnetic fields of such strength and with bulk matter at such densities, so the calculations rely heavily on theory.

Nevertheless, despite the difficulties in developing a complete picture of pulsars, by the summer of 1974 there was consensus on their broad features. They were rapidly rotating neutron stars whose periods were very stable except for a very slow increase with time. It was also clear that the more pulsars we knew about and the more detailed observations we had, the better the chances of unravelling the details.

This is what motivated and guided Hulse and Taylor in their pulsar search. The receiver of the 1000 ft radio telescope at Arecibo was driven so that as the Earth rotated, in one hour, the instrument could observe a strip of sky 10 arcminutes wide by 3 degrees long. At the end of each day's observations the recorded data were fed into a computer, which looked for pulsed signals with a well-defined period. If a candidate set of pulses was found, it had to be distinguished from terrestrial sources of spurious pulsed radio signals, such as radar transmitters and automobile ignition systems. The way to do this was to return later to the portion of the sky to which the telescope was pointing when the candidate signals were received and see if pulses of almost exactly the same period were present. If so, they had a good pulsar candidate that they could then study further, such as by measuring its pulse period to the microsecond accuracy characteristic of other pulsars. If not, forget it.

The day-to-day operation of the program was done by Hulse, while Taylor made periodic trips down from Amherst throughout the summer to see how things were going, or communicated by telephone (this was before the electronic-mail era). On July 2, Hulse was by himself when the instruments recorded a very weak pulsed signal. If the signal had been more than four percent weaker, it would have fallen below the automatic cut-off that had been built into the search routine, corresponding to a signal-to-noise ratio of seven, and would not even have been recorded. Despite its weakness, it was interesting because it had a surprisingly short period, only 0.059 seconds. Only the Crab pulsar had a shorter period. This made it worth a second look, but it was August 25 before Hulse got around to it.

The goal of the August 25 observing session was to try to refine the period of the pulses. If this were a pulsar, its period should be the same to better than a microsecond over several days, because even if it were slowing down as quickly as the Crab is, the result would be a change of less than a microsecond. Now the troubles began. Between the beginning and the end of the two-hour observing run, the computer analyzing the data produced two different periods for the pulses, differing by almost 30 microseconds. Two days later, Hulse tried again, with even worse results. As a result, he had to keep going back to the original discovery page in his lab notebook and cross out and reenter new values for the period. Hulse's reaction was natural: annoyance. Because the signal was so weak, the pulses were not clean and sharp like those from other pulsars, and the computer must be having problems getting a fix on the pulses. Perhaps this source isn't worth the hassle. After all, the summer's observations had already yielded a rich harvest of over 40 pulsars, a significant increase in the world's knowledge, so there were many promising candidates to study. If Hulse had actually adopted this attitude and dumped the candidate, he and Taylor would have been the astronomical goats of the decade. As it turned out, the suspicious Hulse decided to take an even closer look.

During the next several days, Hulse wrote a special computer program designed to get around any problems that the standard program might be having in resolving the pulses. But even with the new program, data taken on September 1 and 2 also showed a change in pulse period, a steady decrease of about 5 microseconds during the two-hour runs. This was much smaller than before, but still larger than it should be, and it was a decrease instead of the expected increase. To blame this still on the instruments or the computer was tempting, but not very satisfying.

But then Hulse spotted something. There was a pattern in the changes of the pulse period! The sequence of decreasing pulse periods on September 2 appeared to be almost a repetition of the sequence of September 1, except it occurred 45 minutes earlier. Hulse was now convinced that the period change was real and not an artifact.

But what was it? Had he discovered some new class of object: a manic-depressive pulsar with periodic highs and lows? Or was there a more natural explanation for this bizarre behavior? The fact that the periods nearly repeated themselves gave Hulse a clue as to an explanation. The source was indeed a well-adjusted pulsar, but it wasn't alone!

The pulsar, Hulse postulated, was in orbit about a companion object, and the variation in the observed pulse period was simply a consequence of the Doppler shift. When the pulsar is approaching us, the observed pulse period is decreased because of the Doppler effect, and when it is receding from us the pulse period is increased. Actually, optical astronomers are very familiar with this phenomenon in ordinary stars. As many as half the stars in our galaxy are in binary systems, and since it is rarely possible to resolve the two stars telescopically, they are identified by the up and down Doppler shifts in the frequencies of the spectral lines of the stars (Fig. 1). Here the pulse period plays the same role as the spectral line in an ordinary star. In most ordinary stellar binary systems, the Doppler shifts of the

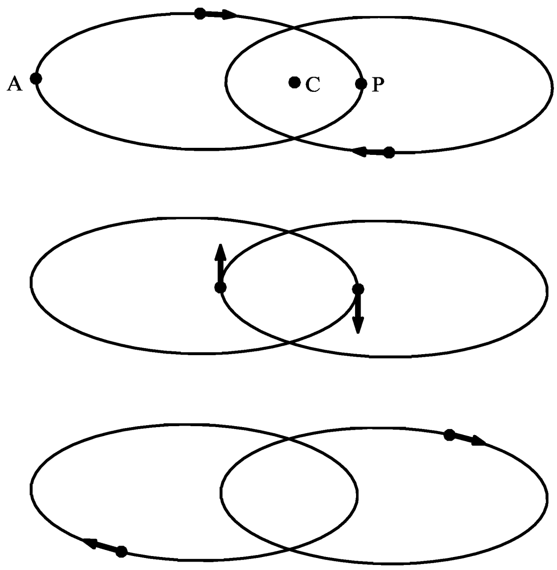

Figure 1. Orbit of a binary system such as the one containing the binary pulsar. Orbit of each body is an ellipse about the center of mass C of the system. Periastron of one body is point P, apastron is point A.

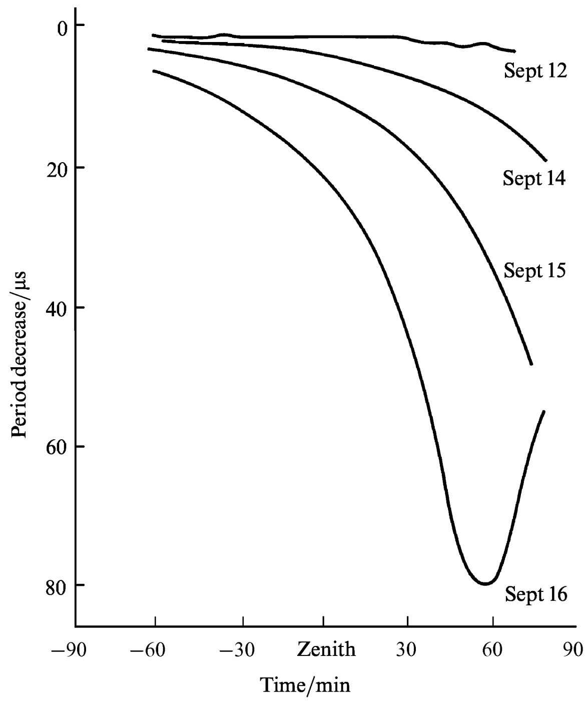

Figure 2. Pulse period changes of the binary pulsar. Data from Hulse's notebook.

spectra of both stars are observed. However, occasionally, one of the stars is too faint to be seen, so astronomers can detect the motion of only one of the stars. Such appeared to be the case here. One of Hulse's problems with this hypothesis was a practical one: he couldn't find any decent books on optical stellar binary systems in the Arecibo library since radio astronomers don't usually concern themselves with such things.

Now, because the Arecibo telescope could only look at the source when it was within one hour on either side of the zenith or overhead direction (thus the two-hour runs), Hulse couldn't just track the source for hours on end, he could only observe it during the same two-hour period each day. But the shifting of the sequence of periods in the September 1 and 2 data meant that the orbital period of the system must not be commensurate with 24 hours, and so each day he could examine a different part of the orbit, if indeed his postulate was right. On Thursday, September 12, he began a series of observations that he hoped would unravel the mystery (Fig. 2).

On September 12, the pulse period stayed almost constant during the entire run. On September 14, the period started from the previous value and decreased by 20 microseconds over the two hours. The next day, September 15, the period started out a little lower and dropped 60 microseconds, and near the end of the run it was falling at the rate of a microsecond per minute. The speed of the pulsar along our line of sight must be varying, first slowly, then rapidly. The binary hypothesis was looking better and better, but Hulse wanted to wait for the 'smoking gun', the clinching piece of evidence. So far the periods had only decreased. But if the pulsar is in orbit, its motion must repeat itself, and therefore he would eventually be able to see a phase of the orbit when the pulse period increases, ultimately returning to its starting value, to continue the cycle.

He didn't have long to wait. The very next day, September 16, the period dropped rapidly by 70 microseconds, and with only about 25 minutes left in the observing run, it suddenly stopped decreasing, and within 20 minutes it had climbed back up by 25 microseconds.

This was all Hulse needed, and he called Taylor in Amherst to break the news. Taylor flew immediately down to Arecibo, and together they tried to complete the solution of this mystery. However, the real excitement was still to come.

The first thing they determined was the orbital period, by finding the shortest interval over which the pattern of pulse readings repeated themselves. The answer was 7.75 hours, so the 45 minute daily shift that Hulse had seen was just the difference between three complete orbits and one Earth day.

The next obvious step was to track the pulse-period variations throughout the orbit to try to determine the velocity of the pulsar as a function of time. This is a standard approach in the study of ordinary binary systems, and a great deal of information can be obtained from it. If we adopt Newtonian gravitation theory for a moment, then we know that the orbit of the pulsar about the center of mass of the binary system is an ellipse with the center of mass as the focus. The orbit of the companion is also an ellipse about this point, but since the companion is unseen, we don't need to consider its orbit directly. The orbit of the pulsar lies in a plane that can have any orientation in the sky. It could lie on the plane of the sky, or we could be looking at the orbit edge on, or its orientation could be somewhere between these extremes. We can forget the first case, because if it were true then the pulsar would never approach us or recede from us and we would not detect any Doppler shifts of its period. We can also forget the second case, because if it were true then at some point the companion would pass in front of the pulsar (an eclipse) and we would lose its signal for a moment. No such loss of the signal was seen anywhere during the eight-hour orbit.

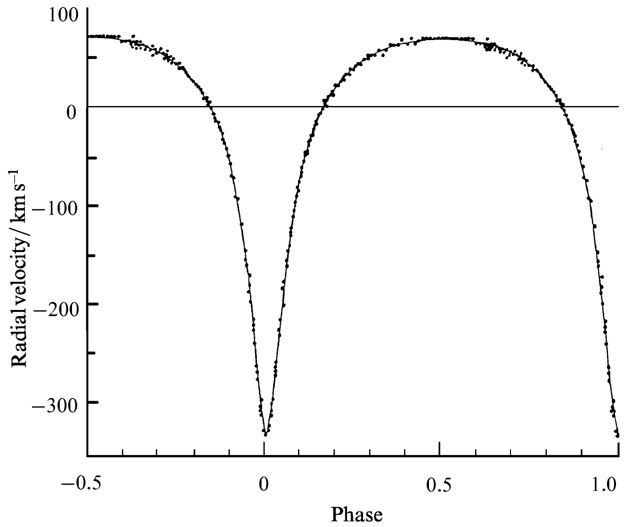

Figure 3. Radial velocity curve for 1.5 cycles of the binary pulsar. Asymmetric shape of curve indicates large eccentricity. Periastron corresponds to sharp dip in curve, apastron to flat portion (courtesy R A Hulse and J H Taylor).

So the orbit must be tilted at some angle relative to the plane of the sky.

That is not all that can be learned from the behavior of the pulsar period. The Doppler shift tells us only the component of the pulsar velocity along our line of sight; it is unaffected by the component of the velocity transverse to our line of sight. If, for example, the orbit were a pure circle, the pattern of Doppler shifts would be sinusoidal and totally unlike the actual pattern observed (Fig. 3). Over a very short period of time (only two hours out of the eight) the Doppler shift went quickly from zero to a large value and back — corresponding to passage through periastron, the point of closest approach of the two stars — while over the remaining six hours, it changed slowly from zero to a smaller value in the opposite sense and back — corresponding to passage through apastron. In fact, the September 16 'smoking gun' observation saw the pulsar pass through periastron. Detailed study of this curve showed that the separation between the two bodies at apastron was four times larger than their separation at periastron, yielding an eccentricity of around 0.6. It also showed that the direction of the periastron was almost perpendicular to our line of sight, since the periastron (the point of most rapid variation in velocity) coincided with the largest Doppler shift (the point where the pulsar has the smallest amount of transverse motion).

At this point, things began to get hot. The actual value of the velocity with which the pulsar was approaching us, as inferred from the decrease in its pulse period, was about 300 km s^{-1}! The velocity of recession was about 75 km s^{-1}. These are high velocities! The speed of the Earth in its orbit about the Sun is only 30 km s^{-1}. Furthermore, if 200 km s^{-1} represents a rough average of the orbital velocity of the pulsar, then the circumference of an orbit that it could trace out in eight hours would be about 6 million km, or about the same as the circumference of the Sun.

When news of this discovery began to spread in late September 1974, it caused a sensation, especially among general relativists. The reasons are as follows. Relativists are always on the lookout for systems in the laboratory or in astronomy where the effects of relativity may be important. They determine this in two ways. First, they calculate $(v/c)^2$, where v is a characteristic speed of an object in the system, and c is the speed of light. The closer this quantity is to unity, the larger are the effects of special relativity, and the happier they are. For the Hulse–Taylor pulsar (200 km s^{-1}) this quantity is about 5×10^{-7}. This is not all that large, except when you compare it to the corresponding value for Mercury, an object where we know relativistic effects are important, for example in the perihelion shift. For Mercury (48 km s^{-1}), this quantity is only 2.5×10^{-8}, a factor of 20 smaller. Secondly, relativists like to calculate GM/Rc^2, where M is a characteristic mass of an object in a system, G is the gravitational constant, and R is a characteristic size (in this case the separation between the bodies). This number is a rough measure of how large the deviations from flat spacetime are near the system, and the bigger it is, the better. For a black hole, it is 0.5. Using two solar masses as the mass of two neutron stars, we find that this number for the pulsar binary system is around 4×10^{-6}, while for Mercury, using the mass of the Sun and the orbital radius of Mercury, it is 2.5×10^{-8}. This made relativists happy indeed, for on the face of it, the orbit of this pulsar was twenty to a hundred times more relativistic than Mercury's orbit. But that wasn't all. Because the orbital period of the pulsar was only eight hours, more than 1000 orbits would occur each year, so that any relativistic effect that built up orbit after orbit, such as the periastron shift (the binary system analogue of the perihelion shift), would build up over 250 times faster than an otherwise equivalent effect for Mercury, which makes only four orbits per year.

It was immediately clear that this new system, called the binary pulsar, or PSR 1913 + 16 after its designation in pulsar catalogues, was a new laboratory for observing general relativistic effects, and it was unique, because it was the first such laboratory outside the solar system. During the Fall of 1974, relativists and astrophysicists swamped the editorial offices of *Astrophysical Journal, Letters* with papers extolling the virtues of this new system and describing all the relativistic effects that could be observed in it. During an eight-week period in early 1975, this journal published seven papers of this kind, in addition to the Hulse–Taylor discovery paper. Between 1975 and 1977, over 40 papers reporting either observational results or theoretical interpretations were published in a variety of astronomical journals, not quite the size of the output over the original pulsar but still a significant cottage industry of research on one object.

Even before Hulse and Taylor's paper on the binary pulsar appeared in print (but too late to stop the presses), Taylor and his colleagues had detected the first of several important relativistic effects, the periastron shift of the orbit. As we have already seen, from the initial observations of these Doppler shifts of the pulsar period it was clear that the periastron line was perpendicular to our line of sight. However, as time went on, the data revealed that it had rotated slightly, by about a third of a degree in one month. During a two-and-a-half month observing program that ended on December 3, 1974, Hulse and Taylor tried to pin down this rotation. Coming up was the Seventh Texas Symposium on Relativistic Astrophysics in Dallas. The data analysis was completed just in time for Taylor to reveal to the audience on December 20 that the rate of periastron advance for the binary pulsar was 4.0 ± 1.5 degrees per year. He would return to the Texas Symposium four years later with an even more impressive announcement.

This periastron advance is about 36,000 times larger than the perihelion advance for Mercury, in keeping with what we expected: a factor 20 to 100 in the raw size of relativistic effects, and a factor of 250 in the number of orbits per year. Is this a triumph for general relativity? It is, but not in the obvious sense. The trouble is, the prediction of general relativity for the periastron advance for a binary system depends on the total mass of the two bodies; the larger the mass, the larger the effect. It also depends on other variables, such as the orbital period and the eccentricity of the orbit, but these are known from the observations. Unfortunately, we do not know the masses of the two bodies with any degree of accuracy. All we know is that they are probably comparable to that of the Sun, in order to produce the observed orbital velocity, but there is enough ambiguity—particularly in the tilt of the orbit with respect to the plane of the sky—to make it impossible to pin the masses down any better from the Doppler shift measurements alone. Well, if we can't test general relativity using the periastron shift measurement, what good is it?

It is actually of tremendous good, because we can turn the tables and use general relativity to weigh the system! If we assume that general relativity is correct, then the predicted periastron shift depends on only one unmeasured variable, the total mass of the two bodies. Therefore, the measured periastron shift tells us what the total mass must be in order for the two values to agree. From the Fall 1974 observations, the inferred total mass was about 2.6 solar masses. Eventually, the periastron shift could be measured so accurately, 4.22663 degrees per year, that the total mass of the system was pinned down to 2.82843 solar masses. This was the triumph for general relativity. Here, for the first time, the theory was used as an active tool in making an astrophysical measurement, in this case the determination of the mass of a system to a few parts in a million.

The relativists' intuition that this system would be a new laboratory for Einstein's theory was confirmed. But there was more to come.

During the first few months of observations of the pulsar, it was realized that this was a very unusual pulsar, over and above its being in a binary system. Once the periodic variations in its observed pulse period were seen to be due to Doppler shifts resulting from its orbital motion these variations could be removed from the data, allowing the observers to examine the intrinsic pulsing of the object as if it were at rest in space. Its intrinsic pulse period was 0.05903 seconds, but if it was slowing down as do other pulsars, it was doing so at an unbelievably low rate. It took almost an entire year of observation to detect any change whatsoever in the pulse period, and when the data were finally good enough to see any change, it turned out to be only a quarter of a nanosecond per year. This was 50 000 times smaller than the rate at which the Crab pulsar's period changed. Clearly, any friction that the spinning neutron star was experiencing was very, very small. At this rate, the pulsar would change its period by only 4 percent in a million years. The steadiness and constancy of this pulsar made it one of the best timepieces the universe has ever seen!

This made it possible for the observers to measure the changes produced in the period by the orbital motion of the pulsar with better and better accuracy. The pulsar was so steady that Taylor and his colleagues could keep track of the radio pulses as they came into the telescope, and even when they had to interrupt the observations for long periods of time—as long as six months—while they returned to their home universities for such mundane duties as teaching, or while the telescope was used for other observing programs, they could return to the telescope after such breaks and pick up the incoming train of pulses, without losing track of a single beep. Eventually, the accuracies with which they could determine the characteristics of the pulsar and the orbit began to boggle the mind: for the intrinsic pulsar period, 0.059029997929613 seconds; for the rate at which the intrinsic pulse period was increasing, 0.272246 nanoseconds per year; for the rate of periastron advance, 4.22663 degrees per year; for the orbital period, 27 906.980895 seconds. Since the pulsar period changes by the quoted amount in the last six digits each year, the measured pulsar period is usually referred to a specific date—in this case, July 7, 1984.

There was more to this accuracy than just an impressive string of significant digits. This accuracy also yielded two further relativistic dividends.

The first of these was another example of 'applied relativity', or relativity as the astrophysicist's friend. Beside the ordinary Doppler shift of the pulsar's period, there are two other phenomena that can affect it, both relativistic in nature. The first is the time dilation of special relativity: since the pulsar is moving around the companion with a high velocity, the pulse period measured by an observer foolish enough to sit on its surface (he would of course be crushed to nuclear density) is shorter than the period observed by us. In other words, from our point of view the pulsar clock slows down because of its velocity. Because the orbital velocity varies during the orbit, from a maximum at periastron to a minimum at apastron, the amount of slowing down will be variable, but will repeat itself each orbit. The second relativistic effect is the gravitational redshift, a consequence of the general relativistic principle of equivalence. The pulsar moves in the gravitational field of its companion, while we the observers are at a very great distance; thus, the period of the pulsar is redshifted, or lengthened. This lengthening of the period is also variable, because the distance between the pulsar and the companion varies from periastron to apastron, and it also repeats itself each orbit. The combined effect of these two phenomena is a periodic up-and-down variation in the observed pulsar period, over and above that produced by the ordinary Doppler shift. But whereas the Doppler shift changed the pulse period by several parts in 10^3, these effects—being relativistic—are much smaller, at several parts in 10^6. It is extremely difficult to measure such a small periodic variation, given the inevitable noise and fluctuations in such sensitive data, but within four years of continual observation and improvement in the methods the effect was found, and the size of the maximum variation was 58 nanoseconds in the pulse period. Again, as with the periastron, this observation does not test anything, because the predicted effect turns out to contain another unknown parameter, namely the relative masses of the two bodies in the system. The periastron shift gives us the total mass, but not the mass of each. Therefore, we can once again be 'applied relativists' and use the measured value of this new effect to determine the relative masses. The result is that the two masses must be very nearly equal, so that if the total mass is 2.828 solar masses, the individual masses must be about 1.42 solar masses for the pulsar and about 1.40 solar masses for the companion, good to about 2 percent. The use and understanding of relativistic effects here played a central role in the first precision determination of the mass of a neutron star.

These results for the masses of the two bodies were also interesting because they were consistent with what astrophysicists thought about the companion to the pulsar. Since it has never been seen directly, either in optical, radio, or X-ray emission, we must use some detective work to guess what it might be. It certainly cannot be an ordinary star like the Sun, because the orbital separation between the pulsar and the companion is only about a solar radius. If the companion were Sun-like, the pulsar would be plowing its way through the companion's outer atmosphere of hot gas, and this would cause severe distortions in the radio pulses that must propagate out of this gas, distortions that are not seen. Therefore, the companion must be much smaller, yet still have one-and-a-half times the mass of the Sun. Such astronomical objects are called 'compact' objects, and astrophysicists know of only three kinds: white dwarfs, neutron stars, and black holes. The currently favored candidate for the companion is another neutron star, based on computer simulations of how this system might have formed from an earlier binary system of two massive stars that then undergo a series of supernova explosions that leave two neutron-star cinders. The fact that both masses turn out to be almost the same is consistent with the observation that in these computer models the central core of the presupernova star tends to have a mass close to 1.4 solar masses. After the outer shell of each star is blown away, the leftover neutron stars each have about this mass. This mass is called the Chandrasekhar mass, after the astrophysicist S Chandrasekhar, who determined in 1930 that this value was the maximum mass possible for a white dwarf (this discovery earned 'Chandra' a share of the Nobel Prize in physics in 1983). Since a presupernova core is similar in many respects to a white dwarf, it is not surprising that this special mass crops up here as well.

So why don't we see the companion? Since the binary pulsar is estimated to be about 16 000 light years away, neither a white-dwarf companion nor a blob of hot gas falling into a black hole would be bright enough to be detectable on Earth. A neutron-star companion would also be much too faint to be seen, unless it, too, were a pulsar. However, there is absolutely no evidence for any pulsed radio waves other than those from the main pulsar, so if the companion is a pulsar its rotating beam must be pointing off in some other direction. Perhaps some distant advanced civilization with its own Hulse and Taylor is watching that pulsar and speculating on the nature of its companion! The currently favored scenario actually suggests that the companion is a dead neutron star, not emitting significant pulsed signals at all. The pulsar that we see actually 'died' once, according to this scenario, but was reborn by being spun up to a very high rotation rate because of friction with the atmosphere of its companion, when its companion was still a normal star. When the companion itself underwent a supernova explosion and core collapse, the resulting pulsar spun down and died on a relatively short timescale, as do most pulsars, but then—lacking a companion with an atmosphere—had no opportunity for rebirth.

But the biggest payoff of the binary pulsar was yet to come. To understand what this payoff was, we must skip back, first to 1916, then to the late 1960s, and finally to Munich, Germany, 1978.

Einstein was not content simply to publish the general theory of relativity and to let matters end there, he continued for several years to study some of the consequences of the theory before turning most of his attention toward his ill-fated search for a unified field theory. One of these consequences was gravitational waves. According to special relativity, the speed of light represents the limiting speed for all interactions. Since general relativity was designed to be compatible with special relativity at some level, you would expect the theory to incorporate such a limiting speed for gravitational interactions, and thus to predict gravitational waves.

In fact, the equations of general relativity did admit gravitational waves as solutions. For example, a dumbbell rotating about an axis passing at right angles through its handle will emit gravitational waves that travel at the speed of light. But Einstein also found that the waves have a very important property: they carry energy away from the rotating dumbbell, just as light waves carry energy away from a light source. He even derived a formula to determine the rate at which energy would be lost from a system, such as a rotating dumbbell, as a consequence of the emission of gravitational waves. As it turned out, the assumptions that he made to simplify the calculation were not completely valid, and he also made a trivial mathematical error that made his answer two times too large, but the basic result was correct.

Einstein's paper on gravitational waves was published in 1916, and that was about all that was heard on the subject for over 40 years. One reason was that the effects associated with gravitational waves were extremely tiny. Another reason was that for a long time there was disagreement over whether the waves were 'real', or whether they were some artifact of the mathematics that would not have observable consequences. But by 1960, two developments resurrected the idea of gravitational radiation. One was the rigorous proof by relativity theorists that gravitational radiation was in fact a physically observable phenomenon, that gravitational waves do carry energy and that a system that emits gravitational waves should lose energy as a result. The second was the decision by Joseph Weber of the University of Maryland to begin to build detectors for gravitational waves from extraterrestrial sources.

By 1974, gravitational radiation was a hot subject, and relativists were dying to find some. Even though Weber had claimed detection of waves as early as 1968, later experiments by other workers had failed to confirm his results, and the general feeling was that gravitational waves had not yet been found. Therefore, when the binary pulsar was discovered, and it was seen to be a new laboratory for relativistic effects, it seemed like a godsend. For if a rotating dumbbell can emit gravitational waves, then so can the rotating binary system, even though the two balls of the dumbbell are held together by a rod, whereas the two stars of the binary system are held together by gravity (in general relativity it doesn't matter what holds them together). The binary pulsar could be used in the search for gravitational waves!

But not in the obvious sense. Because of the binary pulsar's great distance, the gravitational radiation that it emits is so weak by the time it reaches the Earth, that it is undetectable by any detectors in the forseeable future. On the other hand, if the waves are carrying energy away from the system, it must be losing orbital energy. A loss of orbital energy manifests itself in a speed up of the two bodies and a decrease in their orbital separation. The same phenomenon happens, for example, when an Earth satellite loses energy because of friction against the residual air in the upper atmosphere; as it falls toward Earth it goes faster and faster, yet its total energy is declining, being lost in this case to heat. In the case of the binary pulsar, the speeding up combined with the decreasing separation will cause the orbital period to decrease.

Here was a way to detect gravitational radiation, albeit somewhat indirectly, and a number of relativists pointed out this new possibility in the Fall of 1974, soon after the discovery of the binary pulsar. As I mentioned above, the effects of gravitational radiation are exceedingly weak, and this was no exception. The predicted rate at which the 27 000 second orbital period should decrease was only of the order of some tens of microseconds per year. Although this was an exciting possibility, the size of the effect was daunting, and some thought it would take 10 to 15 years of continual observation to detect it. Perhaps by 1990....

Now flash forward four years, to December 1978: the Ninth Texas Symposium on Relativistic Astrophysics, this time in Munich, Germany. Joe Taylor was scheduled to give a talk on the binary pulsar. Rumor had it that he had a big announcement, and only a few insiders and theorists active in the subject of the binary pulsar knew what it was (I knew because I was scheduled to follow Taylor to present the theoretical interpretation of his results). A press conference had been set up for later in the day. In a succinct, 15 minute talk (a longer more detailed lecture was scheduled for the following day), Taylor presented the bottom line: after only four years of data taking and analysis, they had succeeded in detecting a decrease in the orbital period of the binary system, and the amount agreed with the prediction of general relativity, within the observational errors of about 20 percent. This beautiful confirmation of an important prediction of the theory was a fitting way to open 1979, the centenary year of Einstein's birth.

It turned out that the incredible stability of the pulsar clock, together with some elegant and sophisticated techniques developed by Taylor and his team for taking and analyzing the data from the Arecibo telescope, resulted in such improvements in accuracy that they were able to beat by a wide margin the projected timetable of 10 years to see the effect. It was these improvements that at the same time allowed them to measure the effects of the gravitational redshift and time dilation, and thereby measure the mass of the pulsar and of the companion separately. This was important, because the prediction that general relativity makes for the energy loss rate depends on these masses, as well as on other known parameters of the system, so they needed to be known before a definite prediction could be made. By 1991, the data were so accurate that the separate masses were determined to be 1.4411 and 1.3874 solar masses, accurate to 0.05 percent. With these values, general relativity makes a prediction of 75.8 microseconds per year for the orbital period decrease. Using data taken through 1991, Taylor and colleagues recently reported an observed value of 76.0 ± 0.3 microseconds per year (Fig. 4).

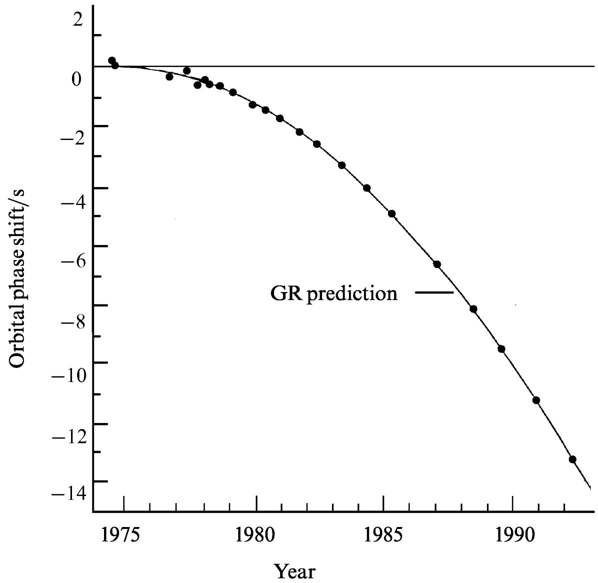

Figure 4. Effect of decreasing orbital period on the phase of the binary pulsar. As the period decreases, the pulsar arrives at periastron progressively earlier. The solid line is the prediction of general relativity with the use of the two measured masses of the stars. The dots are the data points. The experimental error bars are smaller than the dots (courtesy J H Taylor).

Over 40 binary pulsars have now been discovered, and several of them can be used as general relativity laboratories. One of them, PSR 1534 + 12, discovered in 1991 by Alexander Wolszczan, may yield an even more accurate determination of gravitational-radiation damping than did the Hulse–Taylor pulsar. Because these systems contain neutron stars, whose internal relativistic gravitational fields are strong ($GM/Rc^2 \approx 0.2$), the nonlinear nature of gravitation plays a role in the dynamics. Thus, in addition to tests of gravitational radiation, binary pulsar systems have begun to provide important tests of the strong-field nature of gravity, to complement the weak-field tests provided by solar-system experiments.

Joe Taylor and his colleagues continue to study pulsars using high-accuracy timing, making use of upgraded performance of the Arecibo telescope, advanced atomic-clock timekeeping transferred to the telescope by means of the Global Positioning System, and sophisticated data-analysis models. They are studying both binary pulsars and single pulsars (especially those of periods around one millisecond) in order to test fundamental gravitational physics and also to search for gravitational waves from the early universe. In 1980, Joe Taylor moved from the University of Massachusetts to Princeton, where he is now the James S McDonnell Professor of Physics. After receiving his PhD and serving a postdoctoral fellowship, Russell Hulse left the field of radio astronomy for plasma physics and now works in fusion research at Princeton University. On December 10, 1993, Hulse and Taylor received the Nobel Prize in physics from the king of Sweden, for "the discovery of a new type of pulsar, a discovery that has opened up new possibilities for the study of gravitation."

This article is adapted from Chapter 10 of *Was Einstein Right? Putting General Relativity to the Test* 2nd edition, by Clifford M Will (New York: Basic Books, 1993).

Bibliography and suggestions for further reading

Hulse R A, Taylor J H *Astrophys. J. Lett.* **195** L51 (1975)

Taylor J H, Fowler L A, McCulloch P M *Nature* **277** 437 (1979)

Taylor J H, Wolszczan A, Damour T, Weisberg J M *Nature* **355** 132 (1992)

Weisberg J M, Taylor J H, Fowler L A *Scientific American* October, 74 (1981)

Will C M *Theory and Experiment in Gravitational Physics* revised edition (Cambridge: Cambridge University Press, 1993)

Will C M *Was Einstein Right? Putting General Relativity to the Test* 2nd edition (New York: Basic Books, 1993) chapter 10