

Internal Dynamics of Neutron Stars

The central element of many binary x-ray sources is a neutron star. Advances in the theory of dense matter, spurred by precise measurements of changes in spin rates of both rotation-powered and accretion-powered pulsars (see, for example, the section entitled "High-Mass X-Ray Binaries" in the main text) has made it possible to build detailed models describing the properties of these stars. The latest work has led to a dramatic reversal of earlier views.

Theory

From theoretical considerations, neutron stars are thought to consist of several distinct regions (Fig. 1). Immediately below the surface, the matter consists of a solid lattice of more or less ordinary atomic nuclei. As one moves inward to higher densities, however, the protons in the nuclei capture electrons to form neutron-rich nuclei. Above a certain critical density (about 4×10^{11} grams per cubic centimeter), called the *neutron drip point*, it is energetically favorable for some neutrons to be outside nuclei. At these densities the lattice of nuclei is therefore interpenetrated by neutrons, which are be-

lieved to be superfluid. At still higher densities (above about 2.4×10^{14} grams per cubic centimeter—a little below normal nuclear density) the nuclei dissolve completely, leaving a dilute plasma of electrons and superfluid protons in the dense neutron liquid of the core.


In a rotating neutron star the neutron and proton superfluids are expected to behave differently. In both the inner crust and the core the neutron superfluid is thought to rotate by forming arrays of microscopic vortices. These arrays tend to rotate at the same speed as the core. In the inner crust it may be energetically favorable for the neutron vortices to pass through the nuclei of the solid lattice. If this arrangement is sufficiently favorable, the array of vortices will remain fixed in the lattice and will not be able to adjust to changes in the rotation rate of the core. In this case the vortices are said to be *pinned* to the lattice. If the rotational disequilibrium becomes large enough, the constant jiggling of the vortices will cause some to jump from one nucleus to another. This process is called *vortex creep*.

As the rotational disequilibrium builds up, the dynamical forces on a vortex line may exceed the pinning force, causing it to suddenly unpin and move closer to its equilibrium position.

Unlike the neutron superfluid, the proton superfluid is affected by magnetic flux. The magnetic flux in this superfluid is confined to microscopic flux tubes around which proton currents circulate. These flux tubes are much more numerous than the vortices that form due to fluid rotation.

How these various components of a neutron star couple is not well understood, but such coupling determines how the star responds to changes in the rotation rate of the crust. The coupling of the electron and proton fluids in the core to each other and to the solid lattice is thought to be relatively strong, so that one of these components responds to changes in the rotation rates of the others within hundreds of seconds, depending on the rotation rate of the star and whether a strong magnetic field threads the crust and the core. On the other hand, a long-standing view; has been

Fig. 1. The outer crust of a neutron star consists of a solid lattice of nuclei embedded in a sea of relativistic degenerate electrons. The surface of the crust may be solid or liquid, depending on the temperature and the strength of the surface magnetic field. The inner crust is a solid lattice of nuclei embedded in a sea of superfluid neutrons and relativistic electrons. The core is largely a superfluid neutron liquid with a slight admixture of degenerate electrons and superfluid protons. In heavier stars there may also be a distinct inner core that consists of a pion condensate or perhaps a quark soup. Characteristic dimensions are shown. ➤



that the neutron superfluid in the core is only *weakly* coupled to the rest of the star, taking days to years to respond to changes in the rotation rate of the crust. To the extent that this is true, a neutron star can be idealized as consisting of just two components: the crust plus electron and proton fluids in the core, and the neutron superfluid in the core. This idealization has historically been called the *two-component model*. Recent theoretical work and observational data, which we will discuss shortly, has caused astrophysicists to dramatically revise this model.


Imagine forces on the neutron star crust that cause changes in its spin rate (see the section entitled "Pulse Timing" in the main text). Whether the forces are external or internal, whether the crust slows down or speeds up, the stellar interior must eventually adjust. This adjustment can be studied by monitoring the behavior of pulsars after the occurrence of sudden changes in the crust rotation rate. In some rotation-powered pulsars, isolated, relatively large jumps in the rotation rate, called *macroglitches*, have been observed. In both rotation- and accretion-powered pulsars, relatively small fluctuations in the rotation rate have also been seen. The statistical properties of these relatively small fluctuations have been modeled successfully as a series of frequent, small jumps in the rotation rate called *microglitches* (microglitches are thought to be too small to be seen individually with current instruments). The sizes and rates of occurrence of these glitches and the behavior of the rotation rate following them provide tests of models describing the dynamical properties of neutron stars.

Twelve macroglitches have been observed. The largest have occurred in the Vela pulsar (six events) and in three other rotation-powered pulsars (one event each). The smallest macroglitches were one-thousandth the size of the largest Vela macroglitch and occurred in the famous pulsar in the Crab Nebula (two events) and in another rotation-powered pulsar called 0525+21. No isolated events like

these have been seen in accretion-powered pulsars.

Disturbing Results

For a decade, the model used to explain



the behavior of the star following a macroglitch was the two-component model. This model attributes the slowness of the observed recovery of the crust rotation rate, which takes days to months, to the relatively small torque exerted on the crust by the weakly coupled neutron superfluid in the core.

The two-component model appeared to explain the initial data on the post-macroglitch behavior of both the Vela and Crab pulsars. Although model parameters derived from fits to the data were different for the two stars, the parameters were approximately the same for macroglitches in the same star. However, recent detailed analyses of new observations of macroglitches have revealed complex post-glitch behavior not adequately explained by this model.

Further troubling evidence was provided by a detailed study of *microglitches* in the Crab pulsar. Like its response to macroglitches, a star's response to microglitches can be used to probe its internal dynamical properties. This is because the microglitches, which can be described as a random noise proc-



ess, disturb the crust-superfluid system and may be considered an input noise signal applied to this system. The output noise signal, represented by changes in the rotation rate of the crust, is the input noise as filtered by the crust-superfluid system. In other words, the observed power at the analysis frequency f is given by

$$P_{obs}(f) = F(f) P_{in}(f),$$

where $P_{in}(f)$ is the power-density spectrum (Fourier transform) of the forces disturbing the crust, $F(f)$ is the power-transfer function, which reflects the dynamical properties of the neutron star, and $P_{obs}(f)$ is the observed power-density spectrum of fluctuations in the rotation rate of the crust.

Compared to the response times of the system, input forces are expected to be spiky, delta-function-like disturbances. The power-density spectrum of the fluctuating input forces is then a power-law ($P_{in} \propto f^\alpha$ for some constant α) over the analysis frequencies of interest. In fact, one can make $P_{in}(f)$ a constant ($\alpha=0$) by choosing to work with the power-density spectrum of the fluctuations in the right variable (which may be the phase, angular velocity, or angular acceleration of the crust). The shape of the key function $F(f)$ is then given directly by $P_{obs}(f)$. The shape of $F(f)$ can be used to distinguish between neutron-star models, for example, be-

tween a rigid-star model, in which the coupling between the crust and the core is strong, and the two-component model, in which the coupling is weak (Fig. 2). In particular, if a power-density spectrum of the fluctuations in the appropriate variable is flat, this indicates that the star rotates as a rigid body, or, if not, that any internal components are completely decoupled, for disturbances at the analysis frequencies observed.

Unfortunately, at high frequencies noise produced by measurement errors dominates (see the figure in "New Analysis Techniques"), making it difficult to tell whether or not the spectrum remains flat. At low frequencies the spectrum obviously cannot be extended to time scales longer than the observing time. Thus, data from a given experiment can constrain the moments of inertia of components with coupling times only in a certain range. Moreover, power-density estimates always have some uncertainty. Thus, even if the observed spectrum appears to be flat, only an upper bound can be placed on the inertia of components with coupling times in the range studied.

A recent analysis of optical timing data on the Crab pulsar showed that the spectrum of micro-fluctuations in angular acceleration is relatively flat over more than two decades in frequency. This result indicates that the neutron star is responding to microglitches approximately like a rigid body: for a coupling time of 10 days, no more than 70 per cent of the star's moment of inertia can be weakly coupled. These values are sharply inconsistent with the older values derived from fits of the two-component model to macroglitches. These older values implied that 95 per cent of the star's inertia is coupled to the crust with a coupling time of about 10 days. Thus, the two-component model cannot be an adequate description of the full dynamical properties of this neutron star.

Analysis of relatively small-scale fluctuations in the rotation rates of accretion-powered pulsars also shows no evidence of

THEORETICAL POWER-TRANSFER FUNCTIONS




Fig. 2. Logarithmic plots of theoretical power-transfer functions $F(f)$ for three distinct neutron-star models: a rigid star (dashed line), a two-component model (dotted line), and a two-component model in which the vortex array in the superfluid neutron core can be set vibrating (solid line). The shape of the power-transfer function depends on the coupling of the crust to the other components of the neutron star and therefore provides information on the nature of the star's interior.



weakly coupled components. The most complete power-density spectrum is that recently obtained for Vela X-1 (see Fig. 5b in the main text), which covers periods from 0.25 to 2600 days. This spectrum implies that, for coupling times in the range of 1 to 30 days, no more than 85 per cent of the moment of inertia of the star can be weakly coupled.

New Ideas

The inability of the two-component model to account for the response of the Crab pulsar to both macroglitches and

microglitches and the absence of evidence for any weakly coupled component in the microglitch data from the Crab pulsar and Vela X-1 have stimulated theorists to re-examine the view that coupling between

the crust and the neutron superfluid in the core is weak. They found that a previously overlooked quantum liquid effect may account for the inability to detect a weakly coupled component.

As discussed above, the neutron superfluid in the core is expected to rotate by forming arrays of vortices, whereas the proton superfluid in the core is expected to rotate uniformly. Even though the neutrons are superfluid, their motion drags some protons around each neutron vortex, generating a proton supercurrent. (Because the drag coefficient is negative, the proton current actually circulates in the direction opposite to the neutron current.) At the very high proton densities in the core, this induced proton supercurrent generates a magnetic field of 10^{15} gauss (G) near each neutron vortex (even though the mean field in the star generated by this effect is only $10^{-7}/P$ G, where P is the rotation period in seconds). Because of this strong magnetization, the magnetic fields threading the neutron vortices scatter electrons very effectively, causing a strong coupling between the core neutrons and electrons. The resulting short coupling time between the crust and the superfluid neutrons in the core (of the order of $400 P$ seconds) rules out gradual spin-up of these neutrons as the explanation of the long post-macroglitch relaxation times in the Crab and Vela pulsars.

If these theoretical results are correct, the only remaining candidate for a weakly coupled component is the neutron superfluid in the inner crust, where there is no proton fluid and the neutron vortices are therefore unmagnetized. But the moment of inertia of this component is expected to be only about 10^{-2} that of the rest of the star. Thus, a neutron star in which only this component is weakly coupled would behave almost like a rigid body, consistent with the previously puzzling observations of the Crab pulsar and Vela X-1.

What then is the explanation of the macroglitches and the long post-glitch relaxation, which first suggested the idea of weak coupling between the crust and the



core of neutron stars? One possibility builds on the fact that the neutron vortices in the inner crust are expected to be pinned to the lattice of nuclei there. A macroglitch could be a sudden unpinning and movement of many vortices, causing a rapid transfer of angular momentum from the superfluid to the rest of the star and a jump upward of the rotation rate of the crust. If the neutron vortices in the inner crust are dynamically coupled to the crust by vortex creep, then the long relaxation could be explained as the response of this creep process to the macroglitch. This model differs from the two-component model in important ways. Only the neutron superfluid in the inner crust is involved and the superfluid response is fundamentally nonlinear, unlike the linear response given by the frictional coupling of the two-component model.

With relatively few parameters the vortex-unpinning model can explain the post-macroglitch behavior in the Crab and Vela pulsars as well as that in 0525+21. Because post-macroglitch relaxation is thermally activated in this model and hence the relaxation time is proportional to the *internal* temperature of the star, this temperature can be estimated from post-glitch timing observations.

Now if a neutron star is nearly rigid (the theoretical results described above imply that only one per cent of a neutron star is loosely coupled to its crust), the interpreta-

tion of the large pulse-frequency variations seen in accretion-powered pulsars becomes much simpler. Because internal torques can have only a small effect, any substantial fluctuations in the rotation frequency of the crust *must* be due to variations in the external accretion torque, that is, must be due to fluctuations in the torque exerted on the star by material falling onto it.

The theoretical arguments seem persuasive, but it is important to keep in mind that so far we have only a very modest amount of observational evidence. Although analyses of about two dozen rotation-powered pulsars are in progress, detailed microglitch power-density spectra have been published for only one rotation-powered and two accretion-powered pulsars. These spectra only exclude weakly coupled components with moments of inertia greater than 70 to 85 per cent of the star as a whole, for coupling times in the range of 2 to 20 days. Detection of a weakly coupled component with a moment of inertia as small as that expected on the basis of theory (about 1 per cent of that of the star) appears out of reach currently, but the upper bounds on the size of any weakly coupled component can be reduced substantially by more extensive observations using the large-area detectors planned for the near future (see "The Next Generation of Satellites"). Moreover, measurements of a larger number of pulsars are essential before we can be sure that our conclusions about neutron stars are not based on atypical examples. And, of course, there is always the possibility of yet another surprise. ■