Neutron Stars

We now enter the study of neutron stars. Like black holes, neutron stars are very compact objects, so general relativity is important in their description. Unlike black holes, they have surfaces instead of horizons, so they are a lot more complicated than black holes. We'll start with an overall description of neutron stars, then discuss their high densities and strong magnetic fields.

Summary of Neutron Stars

A typical neutron star has a mass of $\sim 1.25 - 2\,M_\odot$ and a radius of $\sim 10 - 15$ km. Spin frequencies are known up to 716 kHz and magnetic fields up to perhaps 10^{15} G or more. Its surface gravity is around $2 - 3 \times 10^{14}$ cm s^{-2}, so mountains of even perfect crystals can't be higher than < 1 mm, meaning that these are the smoothest surfaces in the universe. They have many types of behavior, including pulsing (in radio, IR, optical, UV, X-ray, and gamma-rays, but this is rarely all seen from a single object), glitching, accreting, and possibly gravitational wave emission. They are the best clocks in the universe; the most stable are thousands of times more stable in the short term than the best atomic clocks. Their cores are at several times nuclear density, and may be composed of exotic matter such as quark-gluon plasmas, strange matter, kaon condensates, or other weird stuff. In their interiors they are superconducting and superfluid, with transition temperatures around a billion degrees Kelvin. All these extremes mean that neutron stars are attractive to study for people who want to push the envelope of fundamental theories about gravity, magnetic fields, and high-density matter.

High Densities

Let's start, then, with high densities. An essential new concept that is introduced at high densities is *Fermi energy.* The easiest way to think about this is in terms of the uncertainty principle,

$$\Delta p \Delta x \gtrsim \hbar \; . \qquad (1)$$

Note that, in the spirit of order of magnitude estimates, we are dropping a factor of 2: really $\Delta p \Delta x \geq \hbar/2$, if we want to be precise. If something is localized to a region of size Δx, then its momentum must be at least $\hbar/\Delta x$. That means that in a dense environment, there is a momentum, and hence an energy, associated with the confinement. Therefore, squeezing something increases its total energy, and this Fermi energy acts as a pressure (sometimes called degeneracy pressure). The existence of this energy has a profound role in the structure of white dwarfs, and especially neutron stars. In fact, if degeneracy pressure dominates, then instead of being larger for more massive objects, degenerate stars tend to be *smaller* at higher masses (modulo some complexity related to nuclear physics). In particular, an approximate relation is that $R \sim M^{-1/3}$ for a degenerate star.

First, let's get some basic numbers. If the energy and momentum are low, then the Fermi energy E_F is related to the Fermi momentum $p_F \sim \hbar/\Delta x$ by $E_F \approx p_F^2/2m$, where m is the rest mass of the particle. Since $\Delta x \sim n^{-1/3}$, where n is the number density of the particle, in this nonrelativistic regime $E_F \sim n^{2/3}$. At some point, however, $E_F > mc^2$. Then $E_F \sim p_F c$, so $E_F \sim n^{1/3}$. For electrons, the crossover to relativistic Fermi energy happens at a density $\rho \sim 10^6$ g cm^{-3}, assuming a fully ionized plasma with two nucleons per electron. For protons and neutrons the crossover is at about 6×10^{15} g cm^{-3} (it scales as the particle's mass cubed). The maximum density in neutron stars is no more than $\sim 10^{15}$ g cm^{-3}, so for most of the mass of a neutron star the electrons are highly relativistic but the neutrons and protons are at most mildly relativistic.

Let's now think about what that means. **Ask class:** suppose we have matter in which electrons, protons, and neutrons all have the same number density. For a low density, which has the highest Fermi energy? The electrons, since at low densities the Fermi energy goes like the inverse of the particle mass. **Ask class:** given what we said before, what is the approximate value of the electron Fermi energy when $\rho = 10^6$ g cm^{-3}? That's the relativistic transition, so $E_F \approx m_e c^2 \approx 0.5$ MeV. Then at 10^7 g cm^{-3} the Fermi energy is about 1 MeV, and each factor of 10 doubles the Fermi energy since $E_F \sim n^{1/3}$ in the relativistic regime. What that means is that the energetic "cost" of adding another electron to the system is not just $m_e c^2$, as it would be normally, but is $m_e c^2 + E_F$. It therefore becomes less and less favorable to have electrons around as the density increases.

Now, in free space neutrons are unstable. This is because the sum of the mass-energies of an electron and a proton is about 1.5 MeV less than the mass-energy of a neutron, so it is energetically favorable for a neutron to decay to a proton and an electron (and an electron antineutrino, to balance lepton number). **Ask class:** what happens, though, at high density? If $m_p + m_e + E_F > m_n$, then it is energetically favorable to combine a proton and an electron into a neutron (plus an electron neutrino to balance lepton number). Therefore, at higher densities matter becomes more and more neutron-rich. First, atoms get more neutrons, so you get nuclei such as ^{120}Rb, with 40 protons and 80 neutrons. Then, at about 4×10^{11} g cm^{-3} it becomes favorable to have free neutrons floating around, along with some nuclei (this is called "neutron drip" because the effect is that neutrons drip out of the nuclei). At even higher densities, the matter is essentially a smooth distribution of neutrons plus a $\sim 5 - 10\%$ smattering of protons and electrons. At higher densities yet (here we're talking about nearly 10^{15} g cm^{-3}), the neutron Fermi energy could become high enough that it is favorable to have other particles appear.

It is currently unknown whether such particles will appear, and this is a focus of much present-day research. If they do, it means that the energetic "cost" of going to higher density is less than it would be otherwise, since energy is released by the appearance of other, exotic particles instead of more neutrons. In turn, this means that it is easier to compress the star: squashing it a bit doesn't raise the energy as much as you would have thought. Another way of saying this is that when a density-induced phase transition occurs (here, a transition to other types of particles), the equation of state (which is a relation between thermodynamic variables such as the pressure and energy density) is "soft".

Well, that means that it can't support as much mass. That's because as more mass is added, the star compresses more and more, so its gravitational compression increases. If pressure doesn't increase to compensate, then the star collapses and forms a black hole. What all this means is that by measuring the mass and radius of a neutron star, or by establishing the maximum mass of a neutron star, one gets valuable information about the equation of state, and hence about nuclear physics at very high density. This is just one of many ways in which study of neutron stars has direct implications for microphysics.

Comment: the extra "squishiness" of matter when it is near a density-induced phase transition may also have importance in the early universe. It's been pointed out that when the universe goes from being a quark-gluon plasma to being made of nucleons (at about 10^{-5} s after the Big Bang), this is a density-induced phase transition, so it may be comparatively easy to form black holes then. The electroweak phase transition, which happened earlier, is another possibility. Perhaps one of these transitions led to the formation of so many black holes that they form dark matter; incidentally, because this event happened before big bang nucleosynthesis, no baryon number constraints are violated. This is not the leading model for dark matter, but it is thought-provoking.

Estimating the Masses and Radii of Neutron Stars

From the preceding it is clear that neutron star mass and radius measurements are helpful in constraining the properties of cold high-density matter. The "cold" in that statement means that the temperature is much less than the Fermi temperature, so neutron stars are, especially in their cores, very close to the equilibrium for cold matter.

So how well can we measure neutron star masses and radii? Masses can be straightforward if the star is in a binary. Indeed, the masses of some neutron stars have been measured with the highest precision of any objects outside the Solar System. These are the ones for which post-Keplerian parameters such as pericenter precession, Shapiro delay, and orbital decay due to gravitational radiation can be measured. Two neutron stars in binaries have masses consistent with $\approx 2\ M_\odot$, to $0.04\ M_\odot$ precision. These high masses do place interesting constraints on dense matter, but a lower mass limit alone isn't enough: the maximum mass could be much higher than $2\ M_\odot$, and the radius consistent with this range of masses could run from $10 - 15$ km for a canonical $1.4\ M_\odot$ star.

We'd really like to measure the radius. Unfortunately, existing estimates are completely dominated by systematic errors, to the extent that I don't think we can claim to have any constraint on the radius. As an example (and please see my reviews

http://adsabs.harvard.edu/abs/2013arXiv1312.0029M
and
http://adsabs.harvard.edu/abs/2016EPJA...52...63M

for more details), you might think that you could measure neutron star radii in the way that you can measure the radii of ordinary stars. That is, you could (1) determine the approximate bolometric flux F from the star, (2) determine its distance d, so that (assuming, reasonably, that the radiation is nearly isotropic) the star's luminosity is $L = 4\pi d^2 F$, and (3) fit the star's spectrum to a blackbody, such that the temperature T of the blackbody determines the radius R via $L = 4\pi R^2 \sigma_{\rm SB} T^4$, where $\sigma_{\rm SB}$ is the Stefan-Boltzmann constant.

So what's wrong with that? What's wrong is that radii determined in this manner are of the order of 5 km, which is well under the theoretically reasonable $\sim$ 10 km minimum. The problem with this method turns out to be that although neutron star spectra often look a lot like Planck spectra in the limited range of X-ray energies where we can see them, they are *not* blackbodies. Instead, they have color temperatures (i.e., the temperature you get by fitting a blackbody) that can easily be twice their effective temperatures (which is the temperature you should put into the $L = 4\pi R^2 \sigma_{\rm SB} T^4$ law). To make things more difficult, neutron star spectra do not have identifiable atomic lines; they are just continuum spectra, which means that we don't have any independent way to check our fits.

Now, I'm not giving up on X-ray observations of neutron stars. I'm a member of the NICER satellite team, and also a member of the LOFT collaboration, and I've done a great deal of work that convinces me that there are methods (including fits of the energy-dependent waveforms from rotating hot spots) that have excellent prospects for providing us with radii that are *not* susceptible to these kinds of systematic errors. Nonetheless, for a matter as important as neutron star radii (since it is essential for a key problem in nuclear physics), we really want multiple independent methods, so that we can cross-check our answers.

So how could gravitational waves help? Suppose that we detect the gravitational radiation coming from the coalescence of two neutron stars, and that we compare them with the gravitational waveform we would expect from two black holes of the same mass (even though we don't expect black holes to be as light as neutron stars, we can perform the exercise). When the two objects are well separated from each other, we find that the waveforms are nearly identical. But as they approach within a few stellar radii, tidal effects on the stars cause deviations from the black hole waveform. The larger the radii, for a given mass, the larger the deviations. Thus, it is hoped, the radius can be measured, although it actually turns out that a parameter called the tidal deformability will be what is measured more directly from the waveforms.

Another idea, this time related to mass, was conceived independently by a group led by Chris Fryer, and my group at Maryland. The idea has to do with short gamma-ray bursts. These bursts are believed to be the result of the merger between two neutron stars, or a neutron star and a black hole. Moreover, if the merger is between two neutron stars, it has been argued that the merged remnant has to collapse quickly (in less than 0.1 seconds) to a black hole, because if it does not then there will be a large wind of baryons driven by neutrinos, and this will cause the duration of the burst to exceed what is normally seen.

This means that the total mass of the two neutron stars that merge must be greater than what can be sustained by a rotating neutron star (if it is less, then the rotating remnant is stable and thus does not collapse as needed). Fryer et al., and Lawrence et al. (my group) thus argued that we can obtain an upper limit on the maximum mass of a neutron star, which complements the lower limit of $\approx 2\ M_\odot$ that we get from the two observed systems mentioned earlier. Now, this upper limit depends on the masses of the two neutron stars that merge; if those masses are large, then so is the upper limit, but if the masses are small, the upper limit is small and thus we could have quite a tight constraint on the maximum mass of a slowly rotating neutron star. Indeed, as we showed, if the neutron stars that combine to produce short gamma-ray bursts have masses in the range of the ones we know in our galaxy, the upper limit is just $2.05 - 2.2\ M_\odot$, which is very close to $2\ M_\odot$! But the way to get really reliable constraints is to see a gamma-ray burst and *also* see gravitational waves from the merger, because this way we will know the total mass. I find this to be an exciting prospect for a previously unanticipated constraint on dense matter than will come from gravitational wave detections.

Puzzles Related to Neutron Stars

- What is the equation of state of the matter in the cores of neutron stars? This is the really big question, and as we just said, gravitational wave observations will provide an excellent complement (at least) to electromagnetic observations.
- What are the masses of neutron stars in double neutron star systems? In our Galaxy, the mass range is remarkably narrow: from $1.25\ M_\odot$ to $1.44\ M_\odot$, even though we know other neutron stars with masses up to $2.01\ M_\odot$. The narrow mass range probably means that there is a narrow channel of formation of these objects. Are there significant exceptions we can find? Can we find any neutron stars with masses *lower* than the $\sim 1.2 - 1.25\ M_\odot$ that we think is the minimum that can form from core collapse?
- What are the spins of neutron stars in double neutron star systems? Current wisdom is that rapidly-rotating neutron stars (with periods of a few milliseconds or less) acquire their angular momentum via accretion of matter from a companion. Double neutron

star systems don't have the time for much accretion (at least in the usual picture), because the progenitors of neutron stars are high-mass stars that therefore live just a short time. We therefore expect to see relatively slowly-rotating neutron stars in double neutron star systems, and this expectation is borne out: the shortest period we know in such a system is 23 milliseconds, which might sound fast but it probably > 30 times longer than it could be. It is therefore usually assumed that gravitational waveforms won't be affected much by spin. But are there double neutron star systems with much faster spin?

Really Big Magnetic Fields

In addition to ultrahigh densities, another unique aspect of neutron stars is their magnetic fields. By a factor of $10^6 - 10^7$, neutron stars have the strongest magnetic fields in the known universe. The fields can therefore have extremely important effects on matter in ways not approached anywhere else. Here we'll concentrate on "ordinary" fields of $\sim 10^{12}$ G. For the so-called "magnetars", the field strengths are believed to extend up to $\sim 10^{15}$ G.

Here, by the way, is a place where we can do a good order of magnitude calculation. How strong an average magnetic field would you need to make a significant impact on the structure of the star as a whole? A very rough estimate would involve comparing the mass-energy of the magnetic field to the mass-energy of the star as a whole. The energy density of a magnetic field of strength B is $B^2/8\pi$. This, multiplied by the volume of the star, is the mass-energy in the magnetic field. This needs to be compared with $Mc^2 \approx 2.5 \times 10^{54}$ erg $(M/1.4\ M_\odot)$, where we have scaled to a canonical neutron star mass of 1.4 $M_\odot$. If the volume of the star is around 10^{19} cm^3, then this means that $B^2/8\pi = 2.5 \times 10^{35}$, so $B \approx 2 \times 10^{18}$ G(!!!). Thus field strengths less than "only" 10^{15} G won't have large-scale effects on the stellar structure, although they could have important effects on atomic spectra and energy transport near the surface where the density is much less than average. This, by the way, is the quick answer to why magnetic fields aren't expected to have any detectable influence on the gravitational waveforms of inspiraling neutron stars; it is sometimes suggested that the fields could be amplified significantly by twists caused by the orbit, but in reality the field would balloon out to reduce magnetic energy, and thus the magnetic field can't get amplified (at least before merger) to anything like the magnitude that would be required to have a significant effect on gravitational radiation.

Magnetic fields tend to have a minor impact on the structure of ordinary stars, and even a minor impact on their spectra (fields less than 10^5 G or so are difficult to detect in stars other than the Sun). For neutron stars, however, their effect on the spectrum is dominant, as is their effect on radiation transport properties, which are in turn the most important energy transport mechanisms near the surface of the star.

Let's start the treatment by thinking of a free electron spiraling in a magnetic field. Suppose that it is moving at an angular frequency ω in a circle of radius r, and apply the classical force balance equation. **Ask class:** what is the centrifugal force for an electron, of mass m_e? It's just $m_e\omega^2 r$. **Ask class:** if it is moving at an azimuthal velocity v_ϕ, what is the magnetic force on the electron for a magnetic field B? Simply eBv_ϕ/c. Equating the two, and using $v_\phi = \omega r$, we find the frequency (the cyclotron frequency)

$$\omega_c = \frac{eB}{m_e c} = 11.5\,\mathrm{keV}\,B_{12}\ , \tag{2}$$

where as usual the convention is that $B_{12} = B/10^{12}$ G. Another important quantity is the scale length of the radius of the orbit, which is given by the Bohr-Sommerfeld quantization rule:

$$L_z = m_e\omega r^2 \sim m\hbar \tag{3}$$

where $m = \pm 1, \pm 2, \ldots$ is the azimuthal quantum number. Therefore, the radius of the mth orbital is something like

$$r_m \sim \hat{r}\sqrt{|m|}\ , \tag{4}$$

where $\hat{r} = (\hbar c/eB)^{1/2} = 2.5 \times 10^{-10}\,B_{12}^{-1/2}$ cm is the Landau radius. More exactly, $r_m = \hat{r}\sqrt{2|m|+1}$.

Ask class: we now want to estimate when the orbitals of atoms are affected significantly by the magnetic field. What should we compare? We could compare energies (cyclotron versus Coulomb), forces (magnetic versus Coulomb), or distances (size of atomic orbital versus radius of mth Landau orbital). All three give approximately the same answer. Using forces, equality happens when

$$\frac{Ze^2}{r_m^2} = \frac{ev_m B}{c} = \frac{e\omega_c r_m B}{c}\ , \tag{5}$$

giving the critical field for state m of a hydrogenic atom of nuclear charge Z as

$$B_c = \frac{Z^2}{(2m+1)^3} B_0\ . \tag{6}$$

The critical field for the ground state of hydrogen is $B_0 = m_e^2 c e^3/\hbar^3 = 2.35 \times 10^9$ G, at which point the cyclotron energy is 2×13.6 eV. **Ask class:** given this, are magnetic effects likely to be most important for ground states or excited states? For hydrogen or for heavier atoms? Excited states and low-charge nuclei are easiest to get in the magnetically dominated regime. In fact, lab experiments with high Rydberg level hydrogen ($n \sim 100 - 200!$) have shown some of these effects.

Ask class: consider now a neutron star with a magnetic field $\sim 10^{12}$ G. The energy difference between the ground state and first excited Landau state is $\hbar\omega_c$. At what approximate surface temperature would one expect excited Landau states to exist? At about

$kT = 10$ keV, or $T \sim 10^8$ K. This is much hotter than most neutron star surfaces, which have $T \sim 10^5 - 10^6$ K.

When $B \gg B_c$ and $kT \ll \hbar\omega_c$, atoms become essentially cylindrical (simple when B is very small or very large, but complicated when B is intermediate). In this limit, the Coulomb force dominates only along the magnetic field, and the orbitals are tightly bound across the field. The length across the field for the ground state is therefore $\hat{r}$, and the length along the field is greater, with some length ℓ to be determined. The energy of these orbitals is

$$E \sim \frac{\hbar^2}{2m_e\ell^2} - \frac{Ze^2}{\ell}\ln(\ell/\hat{r}) \ . \tag{7}$$

Minimizing with respect to ℓ gives

$$\ell \sim \left(\frac{a_0/Z\hat{r}}{\ln(a_0/Z\hat{r})}\right)\hat{r} \tag{8}$$

where $a_0 = 5 \times 10^{-9}$ cm is the Bohr radius. For example, when $B = 10^{12}$ G, $a_0 \approx 20\hat{r}$, so $\ell \approx 7\hat{r}$ for hydrogen. The energy is then

$$E \sim Z^2\frac{\hbar^2}{m_e a_0^2}\ln^2\left(a_0/Z\hat{r}\right) \ . \tag{9}$$

The factor before the log is just the usual ground state energy without any magnetic field. For $B = 10^{12}$ G, the multiplying factor in this formula is about 9, predicting about 120 eV for the ground state energy of hydrogen. The real value is about 160 eV.

This is amazing! It means that for a typical NS magnetic field, the ground state energy is *ten times* what it is for no field! This makes a huge difference in many ways.

Superconductivity and Superfluidity

One of the principles that we just encountered is that nature, being lazy, will go for the lowest energy state possible in some circumstance, all else being equal. That's why there is progressive neutronization of matter at higher and higher densities: it's a lower energy state. In that same general spirit, we also can have superconductivity and superfluidity in neutron stars.

The general idea is that if there is an attractive pairing interaction between fermions, then they can couple to form a state with integer spin, and can therefore act like bosons. At a low enough temperature, these "bosons" can form a condensate-like state in which all of the bosons occupy the same quantum state and form a superfluid. If the component fermions are charged, this forms a superconductor. In normal laboratory experience, the pairing is electronic and happens only at very low temperatures (other than the ceramic high T_c superconductors, which do their thing at liquid nitrogen temperatures or a bit above, almost all laboratory superfluid or superconducting phenomena are observed at temperatures less than 20 K). However, in the dense cores of neutron stars, nucleonic pairing can happen.

As with all highly degenerate systems, pairing occurs between states near the Fermi surface (recall that in the cores of NS, both protons and neutrons are degenerate, just not relativistically so). Since there are many more neutrons than protons, neutrons and protons can't pair up easily because their momenta are substantially different. So, consider only $n - n$ and $p - p$ pairings. The first gives a superfluid, and the second gives a superconductor.

Another general principle of phase transitions is that they happen only when it is energetically favorable. In this case, we need to compare the pairing energy Δ with **Ask class:** what other energy? The thermal energy, kT. The pairing energy is *extremely* difficult to calculate from first principles; one reason is that the medium in which the pairings takes place makes a difference (particularly for the outnumbered protons). The value, however, is somewhere around 0.1 MeV. **Ask class:** what does that mean for the approximate transition temperatures? Since 1 eV equates to about 10^4 K, the transition temperature is around 10^9 K (these are the *real* high-temperature superconductors of the universe!). The interior temperatures of neutron stars are expected to be less than this for stars older than a few hundred years at the most, so most of the mass of neutron stars is superconducting and superfluid(!). In the inner crust, between neutron drip (at about 4×10^{11} g cm^{-3}) and nuclear density (about 2×10^{14} g cm^{-3}), the free neutrons are probably paired in the $^1S_0$ state, as in BCS superconductivity (that is, their spins are opposite). At higher densities, calculations suggest they are in the $^3P_2$ state, with aligned spins. Protons are probably coupled in the $^1S_0$ state to form a superconductor.

Ask class: since most of the mass of a neutron star is a superconducting superfluid, what consequences would this have for some of the bulk properties of the star, such as its thermal, rotational, or magnetic properties? A superfluid is irrotational, so any rotation must be quantized in vortices of ordinary matter. However, these vortices are close enough together (about 10^{-2} cm for a 30 Hz rotator like the Crab) that for many purposes one can treat the interior as rigidly rotating. A superfluid also is a perfect thermal conductor, so to an even greater extent than for normal degenerate matter, the interior of a neutron star is isothermal. A superconductor excludes magnetic flux, so any magnetic field in the interior must likewise be quantized in vortices.

Superconductivity and superfluidity, if their effects are observed in NS, could tell us a lot about the pairing and hence inform us about aspects of nuclear physics that are mighty difficult to get from laboratories. This is an extremely indirect process, and too long a story to go into here. Suffice it to say that it has been invoked to explain glitches (sudden but small changes in the spin frequency of pulsars) and the evolution of magnetic fields in neutron stars.