Measurements of neutron star masses and radii

We now have our next "frontiers" lecture. This time, the subject is measurements of neutron star masses and radii, and what they can tell us about very dense matter. This topic is near and dear to my heart: I've worked on it in different ways for many years, and measurement of the masses and radii of a few neutron stars is the prime goal of NASA's Neutron Star Interior Composition ExploreR (NICER; yes, the final R is capitalized) mission. For this lecture I will draw heavily on material that I prepared for two reviews, which you can find at:

http://www.astro.umd.edu/ miller/reprints/miller13c.pdf

and

http://www.astro.umd.edu/ miller/reprints/miller16c.pdf.

The overall perspective

We care about NS mass and radius measurements not just because neutron stars are awesome, but also because if we could get M and R we would learn about aspects of nuclear physics that we can't probe in laboratories. In particular, the matter in the cores of neutron stars is (1) at several times nuclear density, and (2) cold, in the sense that the temperature is much less than the Fermi temperature (which is equal to the Fermi energy divided by Boltzmann's constant k). For such matter there should be far more neutrons than protons, which is a situation we don't encounter in the laboratory. In addition, if other particles (such as hyperons) form the bulk of the mass, their mutual interactions will be very important, and that's tough to judge in labs because we can't produce many of those exotic particles!

To set the stage, I have reproduced two figures from my 2013 review. Figure 1 shows the neutron total energy (rest mass-energy plus Fermi energy) as a function of density, for pure neutron matter. You see that at high densities a few other particles come in with rest mass-energies low enough that they might be energetically favorable, depending on the mutual interaction energies of the exotic particles.

Then Figure 2 shows the gravitational mass versus the circumferential radius for several candidate equations of state, for a nonrotating neutron star.

We will now talk about how neutron star observations can help constrain the properties of dense matter. We will first talk about the relatively well-established masses of some stars, and then talk about the much more poorly understood radii.

Measurements of neutron star masses

This section is mostly taken from Miller and Lamb 2016, EPJA, 52, 63; it is miller16c.pdf in the listing above. Please see the original paper for specific citations to the literature, which we will largely suppress here for readability.

The most precise and reliable measurements of neutron star masses have been made for neutron stars that are in a binary system with another star. This is because (1) via gravity, mass has an effect at a distance; (2) the underlying theory (Newton's laws for Keplerian motion and general relativity for post-Keplerian effects) is well-understood and tested; and (3) many such systems are relatively clean, in the sense that there are no known complicating astrophysical effects that could potentially confuse or bias the dynamical mass measurement. In some cases, measurement of post-Keplerian effects such as pericenter precession, the Shapiro delay, and orbital decay due to gravitational radiation can overdetermine the properties of the system, providing a test of general relativity.

As we learned in Lecture 11, if the orbital period $P_{\rm orb}$ of a binary system can be determined and the periodic changes in the line-of-sight velocity K_1 of one of the stars in the system can be measured, then to Keplerian order one can construct the mass function $f_1(M_1, M_2)$, which provides a lower limit on the mass M_2 of the *other* star. For a circular orbit,

$$M_2 \geq f_1(M_1, M_2) \equiv \frac{M_2^3 \sin^3 i}{(M_1+M_2)^2} = \frac{K_1^3 P_{\rm orb}}{2\pi G}\ . \tag{1}$$

Here M_1 is the mass of the star whose velocity is measured and i is the inclination of the orbit; $i = 0$ means that we are viewing the orbit face-on, whereas $i = 90^\circ$ means that we are viewing the orbit edge-on. $M_1 = 0$ and $i = 90^\circ$ would yield $M_2 = f_1$; because M_1 is > 0, M_2 must be $> f_1$, even if $i = 90^\circ$.

To uniquely determine the masses of both stars in a binary system with a known mass function, at least two additional properties of the system must be measured. The ideal systems for such measurements are double neutron star binaries, because both objects in the binary are effectively point masses relative to their separation (a typical separation is $\sim 10^{11}$ cm $\gg R \approx 10^6$ cm). If periodic pulses can be detected from at least one of the neutron stars, the high precision of pulsar timing can be brought to bear, and at least some of the post-Keplerian parameters mentioned above can be measured and used to remove degeneracies.

The two highest neutron star masses that have been precisely measured were determined in different ways. The mass of the pulsar PSR J1614−2230, which has a half-solar-mass white dwarf companion, was determined by measuring the mass function of the system and the Shapiro delay using radio observations of the pulsar (Demorest et al. 2010, Nature, 467, 1081). The Shapiro delay is a relativistic effect that causes the light-travel time through the gravitational well of a star to be greater than in the Newtonian limit and to vary periodically with the orbital phase of the system relative to our line of sight to the system. The magnitude of the delay depends on the mass of the pulsar's companion, while its variation depends on the inclination of the system relative to our line of sight (e.g., if the system is face-on to us, the delay has no orbital-phase dependence). Importantly, the Shapiro delay does not depend on the *nature* of the companion. It is therefore irrelevant whether the pulsar's companion is a neutron star, a white dwarf, or a main sequence star. Measuring the Shapiro delay determines the two additional system parameters needed to obtain a unique solution for the masses of both stars in the system. The estimated mass of PSR J1614−2230 is $1.928 \pm 0.017\ M_\odot$.

PSR J0348+0432 has a white dwarf companion with observable atmospheric spectral lines (Antoniadis et al. 2013, Science, 340, 448). The periodic variation of the energies of these lines yields a second mass function, while the measured gravitational redshift of the lines can be used to determine the white dwarf mass, closing the system of equations. The estimated mass of PSR J0348+0432 is $2.01 \pm 0.04\ M_\odot$.

As pointed out to me by Scott Ransom, the improved timing techniques that have been developed for pulsar timing arrays are yielding more precise masses for many stars, so there is hope that neutron stars with even higher masses will eventually be discovered. Some higher masses *have* been reported, up to 2.7 $M_\odot$, but these rely on modeling of complex systems, rather than on the straightforward dynamics that is used for the systems mentioned above.

Although the high masses discovered recently have helped greatly in our understanding of dense matter, a glance at Figure 2 shows that we have a lot of work to do: not all of the equations of state represented in that figure would be considered equally likely by nuclear physicists, but although all of the equations of state can sustain a 2 $M_\odot$ neutron star, the radii at that mass range from 10 to 17 kilometers! Clearly, precise and reliable mass measurements would be highly desirable.

Measurements of neutron star radii

Neutron stars are small and distant, so there is no prospect of angularly resolving them and thus obtaining their sizes if we know their distances. Thus we have to use a more indirect method.

One natural possibility is to measure neutron star sizes the way that we measure the radii of ordinary stars. Basically, we (1) assume that the star's spectrum is more or less a blackbody, (2) assume that the whole photosphere of the star radiates uniformly and therefore isotropically, (3) determine the temperature T from the spectrum, (4) measure the flux F, and (5) measure the distance d of the star from us (e.g., by using parallax). Then the luminosity is $L = 4\pi d^2 F$ and it is also $L = \sigma_{\rm SB} 4\pi R^2 T^4$, where $\sigma_{\rm SB}$ is the Stefan-Boltzmann constant. This gives us R. In a few cases where we can angularly resolve the star (e.g., the Sun or Betelguese), and in a growing number of cases where we can get the radius via asteroseismological modeling (largely from Kepler data; this is one of the great scientific triumphs of the mission), those radii agree well with our blackbody modeling.

If we try the same thing with neutron stars, we get absurd answers: often less than 5 kilometers. What's going on?

The problem is that all current methods for determining the radii of neutron stars using their X-ray fluxes and spectra are subject to astrophysical effects that can confuse or bias the radius measurement. Furthermore, in most cases the data are not yet precise enough to determine whether the model being used correctly describes the data. It is therefore possible that a model may provide a statistically good fit and an apparently tight radius constraint but a value for the radius that is strongly biased relative to the true value. In the case of the naive spectral modeling approach above, the problem is that although the X-ray spectral *shape* for a neutron star is close to a Planck function, it is *not* a blackbody; the opacity is dominated by scattering, which means that the surface is an inefficient emitter. Thus to get out a given amount of flux, the temperature has to go up, and since $R \propto T^{-2}$ in the blackbody formula above, we'd infer a radius that is too low.

But now you might think there is a simple solution. Surely we can just do a better job of modeling the spectra, and then use the same principle as above but with the right spectra?

Not so fast! People *have* tried this, of course, and there are various papers that give the resulting radii. For the reason stated above, those radii are often given with high *precision*, i.e., if the model assumptions are all correct then we know the radius quite well. But the *accuracy* is a big question. Among the potential problems are:

1. We might *not* know the spectrum that well. For example, for non-accreting neutron stars there is a question about whether the spectrum is dominated by hydrogen or helium. You might think hydrogen is a given, because the NS surface gravity is so strong that the lightest element present will float to the top, but there is evidence from some young neutron stars (such as Cas A) that heavier elements can dominate on occasion, and a suggestion dating from the 1960s that given enough time the hydrogen might be slowly fused into helium, which could deplete the hydrogen at the photosphere. Moreover, it makes a big difference: helium atmospheres often give 50%(!) larger radii than hydrogen atmospheres, with an equally good fit to the data.

2. Another assumption is that the full surface emits uniformly. This is a crucial assumption whether the star being modeled is a non-accreting neutron star, or one that is undergoing X-ray bursts. But bursts often have oscillations in their intensity that in-

dicate that there is a hot spot on the surface, so clearly the emission isn't completely uniform. Even when there are no detectable oscillations, there can be nonuniformity; we found that in one particular burst source, which has no oscillations, the fraction of the surface that emits strongly decreases by $\sim 20\%$ during the burst! If nonuniform emission is not included in the model, then the neutron star radius will be underestimated.

There is not solid agreement among the groups that do this type of analysis. Some groups typically find $R = 12 - 13$ km, whereas other groups find $R = 10 - 11$ km. Radii of $10 - 11$ km are difficult to reconcile with the high masses that have been seen. In addition, when the details of the burst analyses that give $R = 10 - 11$ km are examined closely, it becomes clear that the spectral models that are being used are explicitly inconsistent with the data. This isn't too surprising; for example, burst sources have accretion disks, which can contribute unmodeled flux to the spectra, but it does mean that these data can't be used to measure neutron star radii.

My great hope for X-ray measurements is that another method will be more successful. This will be the focus of our analysis of *NICER* data. The idea is that if a neutron star has a hotter region ("hot spot") on its surface and gas in the hot spot rotates around the star, a distant observer will see an energy-dependent waveform whose properties are affected by special relativistic effects such as Doppler shifts and aberration, and general relativistic effects such as light deflection. If you know the rotation frequency of the star, then to lowest order you could get the radius from the linear speed on the surface (which feeds into the special relativistic effects), and the mass to radius radio from the fraction of the surface that is visible (which results from general relativistic light deflection).

In reality, there are many degeneracies to consider. For example, suppose that you see a weak modulation of the intensity from a hot spot on a rotating neutron star. Why is it weak? Maybe the spot is very small, so it doesn't contribute much flux. Maybe the spot is close to the rotational axis, so its flux isn't modulated much (a circular spot *on* the rotational axis wouldn't give any modulation at all). Maybe you, the observer, are close to the rotational pole. Maybe the star is quite compact, so that light deflection smears out the spot.

Fortunately the degeneracies aren't total, which means in turn that sufficiently sophisticated analysis of sufficiently informative data can get you good measurements of the mass and radius.

But what encourages me the most is that unlike the previous methods described, this one seems far more immune to systematic errors. In various papers (particularly Lo et al. 2013, ApJ, 776, 19 and Miller and Lamb 2015, ApJ, 808, 31) my colleagues and I have, among other things, generated synthetic data with different assumptions about the spot shape, beaming pattern, spot temperature distribution, and so on, than we use in our standard analysis.

Thus far, we find that if there would be a bias in the inference of M and R (i.e., if we would think that the real values of M and R are strongly excluded), then the actual statistical fit would be terrible. Thus we would be warned. This is unlike the previous methods described, in which you can have a good statistical fit *and* be very biased.

Future methods

For the future, most attention has been focused on the prospects that gravitational waves from coalescing compact binaries will yield mass and radius constraints that are entirely independent of the constraints derived from electromagnetic observations.

The waveforms from the inspiral and merger of two neutron stars or a neutron star and a black hole bear the imprint of the tidal interactions of the stars. Although many high-precision numerical simulations are still needed, early indications are that analytical models of the tidally-influenced waveform are sufficiently accurate to perform reliable parameter estimation. The detection of several to tens of such events may allow discrimination between soft, medium, and hard equations of state, although it could be important to have good prior knowledge of the distribution of neutron star masses.

It may be possible to place useful constraints on the maximum mass of neutron stars by combining gravitational-wave and electromagnetic observations of short gamma-ray bursts. These bursts are thought to be produced by the merger of two neutron stars or a neutron star and a black hole. If two neutron stars merge, then it has been argued that the merged remnant must collapse within $\lesssim 0.1$ s, otherwise the baryonic wind driven by neutrinos will delay and lengthen the bursts beyond the few tenths of a second duration that is observed. This line of argument implies that a successful burst requires that the total baryonic mass of the two original neutron stars exceed the maximum that can be sustained by a uniformly rotating star. In Lawrence et al. 2015 (ApJ, 808, 186; see also Fryer et al. 2015, ApJ, 812, 24) we find that if the masses of neutron stars that merge to produce short gamma-ray bursts are comparable to those of the neutron stars we see in our Galaxy, then the maximum mass of a nonrotating neutron star is ~ 2.05–$2.2\ M_\odot$, which is quite close to the $2.01\ M_\odot$ maximum mass currently observed. Obtaining more reliable constraints will require identification of individual short bursts with specific gravitational-wave events. Because the gamma rays from these bursts are tightly beamed and are therefore seen by only a small fraction of observers, such associations are likely to require signatures in other electromagnetic bands that can be seen over a much broader solid angle (e.g., optical to infrared ratios in kilonovae or possibly scattered X-rays.

Overall, neutron stars are not giving up their secrets easily! We have hope for the future, but there is a lot of work left to do.

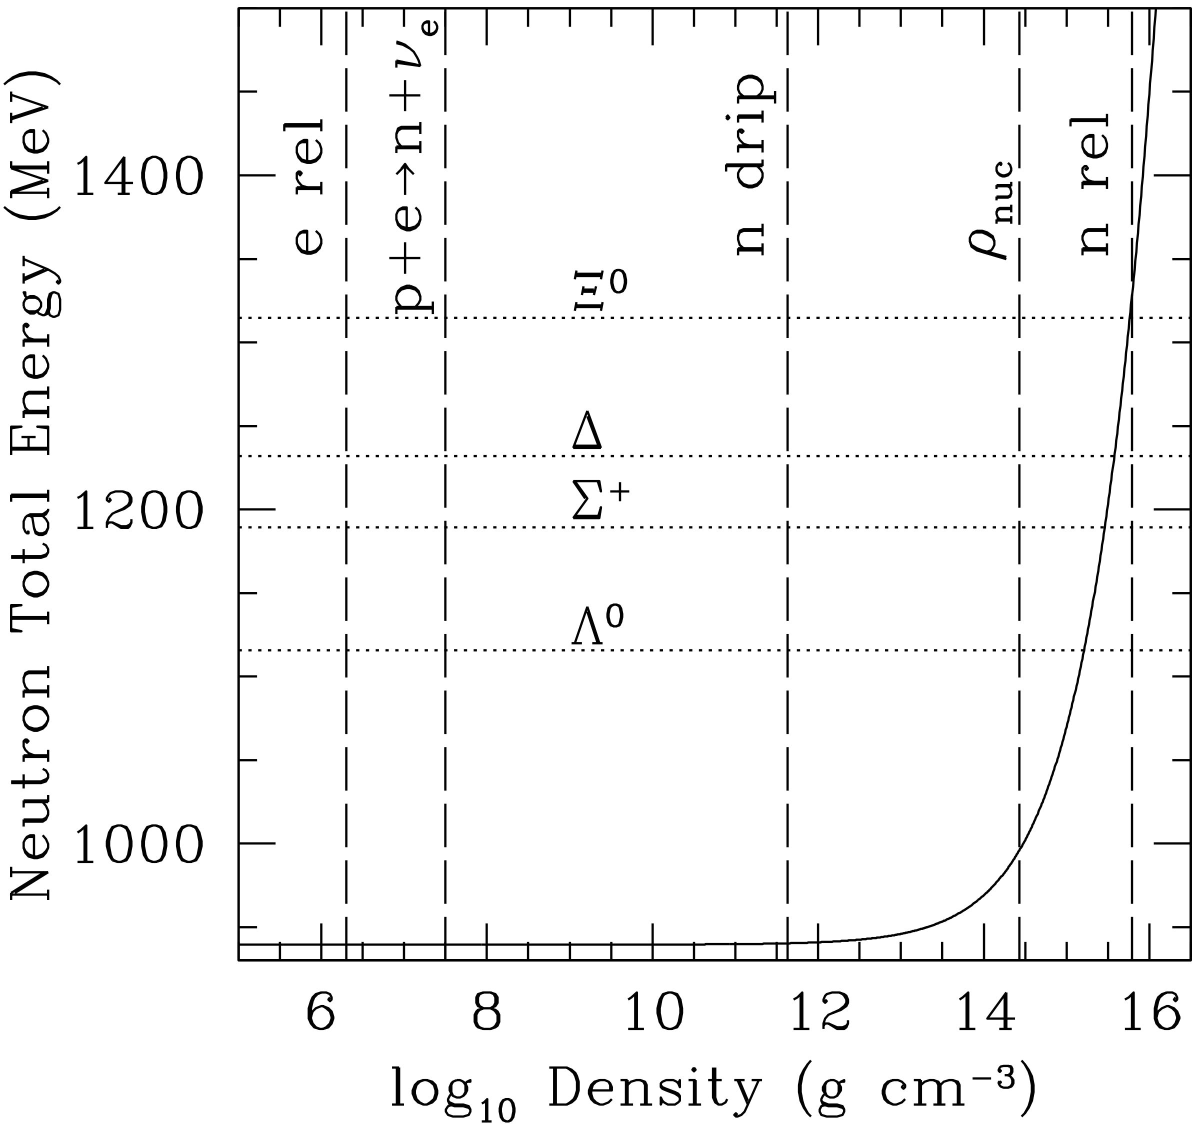

Fig. 1.— Total energy per free neutron versus mass density (solid line). Above $\sim 10^{13}$ g cm^{-3} the Fermi energy starts to contribute palpably to the total, and above $\sim 10^{15}$ g cm^{-3} the total energy can exceed the rest mass energy of particles such as Λ^0, Σ^+, Δ, and Ξ^0 (marked by horizontal dotted lines). Interactions between these particles can change the threshold density. The central densities of realistic neutron stars range from $\sim 5 \times 10^{14}$ g cm^{-3} to $\sim$ few $\times\, 10^{15}$ g cm^{-3}, so some of these exotic particles may indeed be energetically favorable. Also marked are the densities at which free electrons become relativistic; where those electrons have enough total energy to make $p + e^- \rightarrow n + \nu_e$ possible; where free neutrons can exist stably (i.e., at neutron drip); nuclear saturation density $\rho_{\rm nuc}$; and where free neutrons have a Fermi energy equal to their rest-mass energy. To calculate the neutron Fermi energy we assume that all the mass is in free neutrons; in reality at least a few percent of the mass is in protons and other particles, and below $\rho_{\rm nuc}$ a significant fraction of mass is in nuclei. Original figure and caption from Miller 2013, arXiv:1312.0029.

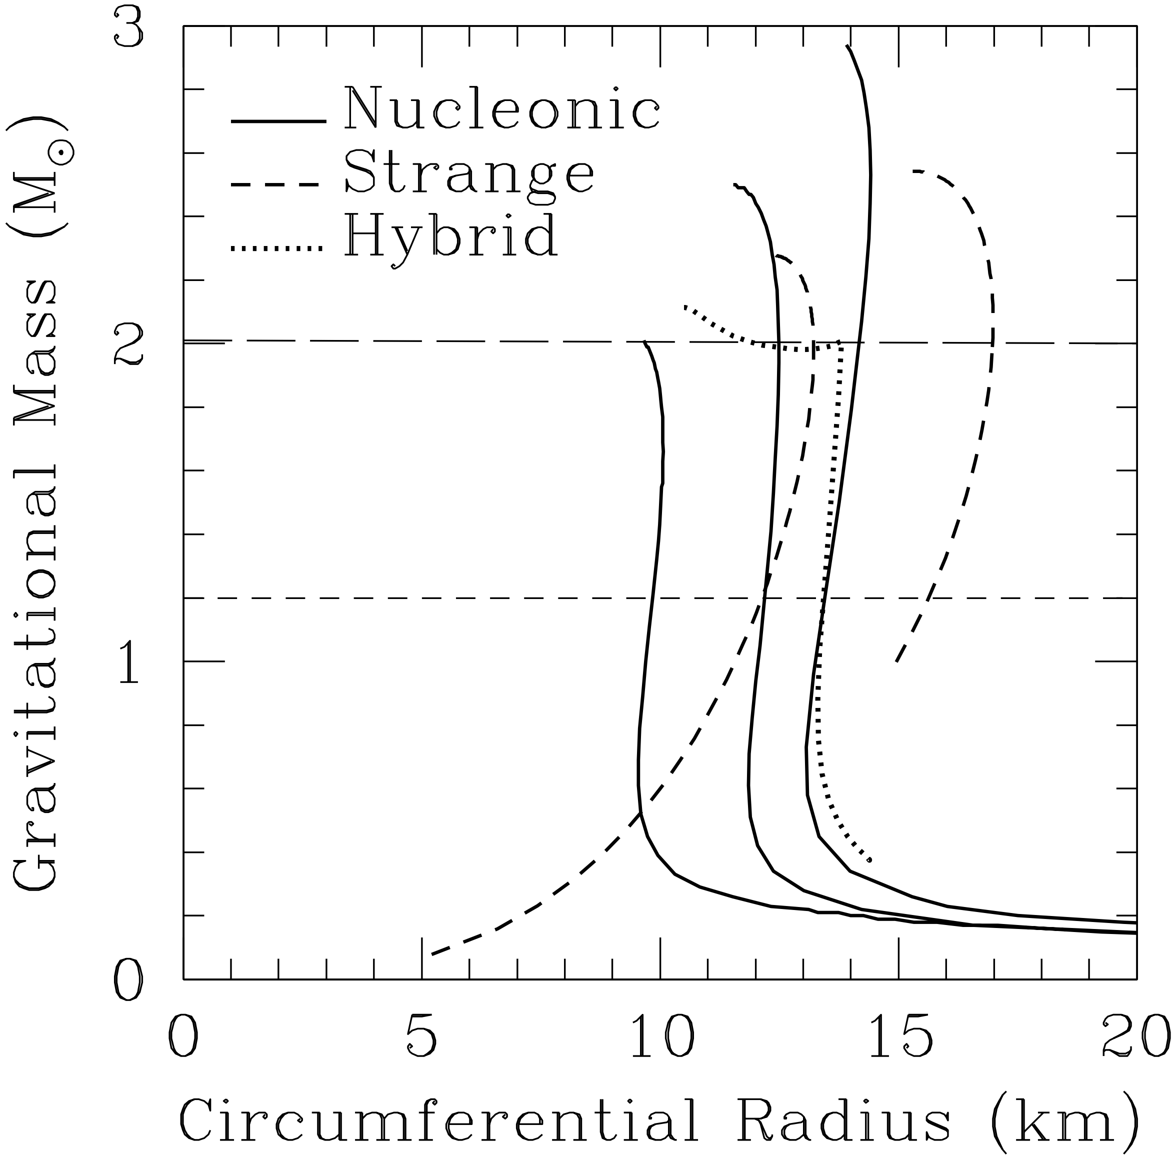

Fig. 2.— Mass versus radius for nonrotating stars constructed using several different high-density equations of state. Rotation changes the radius to second order in the spin rate, but the corrections are minor for known neutron stars. The solid curves include only nucleonic degrees of freedom (these are the mass-radius relations for the soft, medium, and hard equations of state from Hebeler et al. 2013, ApJ, 773, 11), the short dashed lines assume bare strange matter (Kurkela et al. 2010, PRD, 81, 5021), and the dotted curve uses a hybrid quark equation of state with a phase transition (Blaschke et al. 2013, arXiv:1310.3803). The horizontal dashed line at 1.2 $M_\odot$ represents approximately the minimum gravitational mass for a neutron star in current formation scenarios, whereas the horizontal dashed line at 2.01 $M_\odot$ shows the highest precisely measured gravitational mass for a neutron star. Original figure and caption from Miller 2013, arXiv:1312.0029.