X-ray bursts

Last time we talked about one of the major differences between NS and BH: NS have strong magnetic fields. That means that "hot spots" can be produced near the magnetic poles, leading to pulsations (this general statement is true for both rotation-powered and accretion-powered pulsars). Now we'll talk about one manifestation of another crucial difference: NS have surfaces, unlike black holes.

So far we've talked about two sources of energy associated with neutron stars: rotation and accretion. Suppose we have a neutron star that is accreting hydrogen and helium from a companion. **Ask class:** what is another potential source of energy, given that the surface temperature is typically 10^7 K? Since the accreted matter is hydrogen and helium, at such high temperatures nuclear fusion may progress. Deeper down, where the temperature and density are even higher, fusion will definitely happen. Therefore, nuclear burning is another source of energy.

Ask class: if nuclear burning happens continuously during accretion, would it be visible as another source of energy? No. The reason is that nuclear burning is much less efficient than accretion onto a neutron star. Conversion of hydrogen into helium releases a fraction $\approx 7\times10^{-3}$ of the mass-energy, whereas accretion onto the stellar surface releases ≈ 0.2 of the mass-energy, and the uncertainty in the accretion efficiency is a lot more than 7×10^{-3}! Therefore, if burning happens continuously, it will never be noticed.

Moving now to the observational story, starting in 1974 there were a number of detections of bursts of X-rays from different NS LMXBs. The properties varied, but typically the burst would last about 10 seconds, with a rise time that was often less than one second. The time between bursts was anywhere between hours and weeks, depending on the source. In addition, if one defined the parameter α as the ratio of the time-integrated energy in the persistent emission to the time-integrated energy in bursts, the ratio was $\alpha\sim 20-300$. **Ask class:** what could cause this? The total energy release in the bursts is consistent with nuclear burning. It was therefore suggested that here was an example of *unstable* nuclear burning; somehow, the fuel is stored up until some critical moment, then it all burns at once. If the burning happens rapidly, then at least superficially both the rise time and the typical duration can be explained by the radiative transfer time necessary to get, respectively, the first photons out and the last photons out (naturally, as with so much else that we study, a deeper investigation makes this explanation seem less clear). Some bursts are longer: up to thousands of seconds. This is also consistent with the nuclear flash interpretation, if one assumes that it is carbon that is being burned very deep in the star.

For relatively dim persistent sources, the luminosity can go up by a factor of ~ 100 during a burst. The luminosity is, however, limited by the Eddington luminosity. If the luminosity is much higher than that, it creates a wind of matter that streams out and obscures the burst. This is the explanation (first proposed informally by Fred Lamb) for why it is that in high energy photon bins many bursts seem to be "double-peaked". That is, the flux rises, but when it gets high enough to eject matter, the effective radiative surface area increases dramatically (because the surface of optical depth unity is now at a much larger radius). As a result, the emission is much cooler, so although the total luminosity is still high, in the higher energies it decreases. When the luminosity goes well below Eddington, the atmosphere settles down and the high-energy emission increases.

The preceding discussion indicates that a wind of matter can be ejected during a burst. **Ask class:** considering the bulk of the matter involved in the burning, when there is a flash do you expect most of the matter to be ejected to infinity or to be bound fairly closely to the surface? The majority of the matter stays close to the surface, because the burst does not release enough energy to eject more than a small fraction to infinity (this is one consequence of the gravitational binding energy being much larger than the nuclear binding energy).

Ask class: so what would one expect from the same situation on a white dwarf, where there is a nuclear burst? In that case, the radius is 1000x larger than a NS, so the ratio is reversed: nuclear energy dominates over gravitational energy. Therefore, in such a burst it is energetically possible that *more* than the accreted matter is ejected! This is what is thought to happen in classical novae. When the burning starts, the matter is lifted away from the surface. This decreases both the density and the temperature, so the rate of burning drops dramatically. This also means that in novae, unlike in X-ray bursts, burning is thought to be very incomplete. In X-ray bursts, most of the matter is believed to burn to completion.

Instability and different burning regimes

Let's now look more closely at the instability itself. **Ask class:** in general, what does it mean that a system is unstable? It means that a slight perturbation will cause the system to change its state dramatically. **Ask class:**, so what does it mean for the nuclear burning to be unstable? It means that if there is a slight perturbation, the rate of nuclear burning will change quickly. In particular, if for some reason the temperature goes up, then for instability the nuclear burning rate must rise by enough that the temperature goes up further, leading to a runaway process.

This means, first of all, that the rate of nuclear burning must be sensitive to temperature. This is not always the case. Generically, nuclear burning requires that the wavefunctions of nuclei overlap enough for a reaction to happen. In normal thermonuclear burning (such as in stars), this happens because there is a small fraction of nuclei with enough energy to penetrate the Coulomb barrier. Let's treat this classically first. Specifically, imagine that there is some barrier energy E_b that must be overcome for the reaction to occur. If the temperature is

T, then a fraction $\sim \exp(-E_b/kT)$ of the nuclei will have the required energy. If $kT \ll E_b$, this reaction is therefore extremely sensitive to the temperature, since $\exp(-E_b/kT)$ goes up dramatically when T increases. In reality, the process is quantum mechanical rather than classical, and as a result the actual reaction rate is dominated by quantum tunneling. In a very rough way, the effect of tunneling is to lower the barrier energy (that is, a nucleus can react with a much lower energy than one would have thought classically). That's why nuclear reactions can occur in the center of the Sun ($T \approx 2 \times 10^7$ K=2 keV) even though the height of the classical barrier is $E_b \sim 1$ MeV. Qualitatively, then, **Ask class:** would one generally expect the temperature sensitivity to be greater for high or for low temperatures? It is greater for low temperatures, since the reaction rate is exponential. One way people often characterize the temperature sensitivity is by defining a "temperature exponent" ν, with the meaning that at a particular temperature T the reaction rate goes like $R \propto T^{\nu(T)}$. For example, for the CNO cycle of hydrogen burning, the temperature exponent is $\nu \approx 51 T_6^{-1/3}$.

Therefore, except at extremely high temperatures, the nuclear reaction rate goes up with increasing temperature. But how much does it have to go up? The criterion is that there is a *net* increase in temperature. However, remember that if the temperature goes up, so does the cooling rate. Therefore, the energy generation by nuclear reactions must increase faster than cooling. The cooling rate increase like T^4, so for instability in this regime it is necessary that $\nu > 4$. For CNO burning, this means $T < 2 \times 10^9$ K. Above that temperature, burning proceeds steadily; below it, in the thermonuclear regime, burning is unstable.

There is, however, another nuclear burning regime, this time *not* encountered in main sequence stars. We said earlier that the fundamental requirement for burning is that the wavefunction of nuclei overlap. This can happen at arbitrarily low temperature if the density is high enough. This is called "pycnonuclear" fusion, to distinguish it from thermonuclear fusion. **Ask class:** if this form of fusion dominates, will it be very sensitive to temperature? No, because it depends on density instead, so an increase in temperature will do little. Therefore, if the burning is in the pycnonuclear regime, burning is stable and there will be no burst.

One can therefore map out regions in (T, ρ) space for a given composition in which burning is stable or unstable. If the fuel reaches an unstable point, it will produce a burst. Global mapping of such stable and unstable regions has been done by Fushiki and Lamb (1987), and basically similar conclusions were reached by Narayan and Heyl (2003). They follow matter of some composition (usually pure helium or mixed hydrogen-helium) as it settles and is progressively buried. The temperature and density both go up, due to energy release by the settling and by burning. In addition, some of the nuclear fuel is consumed, so there is less to burn. They find that, typically, if the accretion rate is close to Eddington then burning is stable, so there will be few bursts and if there are any they'll be weak.

Quantitatively, ignition typically happens when the column depth of fuel is in the 10^{8-10} g cm^{-2} range. This is only a few tens of meters thick. That's extremely thin compared to the 10 km radius of the star, so to a large extent one can consider this to be a plane-parallel layer.

Propagation of nuclear burning

The vast majority of analyses prior to about 1995 either more or less stopped at this point, or did numerical simulations of bursts in one dimension (radial). However, observational developments since then have changed this dramatically.

The specific development was the discovery with the *Rossi* X-ray Timing Explorer (RXTE) of brightness oscillations during X-ray bursts. Prior to the launch of RXTE in late 1995, no X-ray instrument had the time resolution necessary to detect pulsations with frequencies exceeding about 500 Hz. RXTE, however, was great at such things, and could easily detect oscillations up to 4 kHz with high signal to noise because of its large collecting area. This led to discoveries of two major new phenomena. One is the so-called kilohertz quasiperiodic brightness oscillations, which we won't talk about here. The other is burst oscillations.

There are now more than 20 bursting sources in which oscillations have been detected. Specifically, what is seen is a variation in brightness, which for a given source is essentially always at the same frequency, but which has different frequencies for different sources. The frequency of variation is typically in the $\sim 200 - 600$ Hz range. The consistency of the frequency of these oscillations (for a given source), plus the fact that they are extremely sharp (not perfectly coherent, but $\nu/\Delta\nu > 1000$ in many cases), has led to an almost universally accepted model: once again, we're seeing emission from hot spots that is modulated by rotation. This is similar to rotation-powered and accretion-powered pulsars, leading some people to dub these “nuclear-powered pulsars”.

But wait! If there are hot spots, that means that the nuclear burning is nonuniform. That means that ignition has to happen at one spot first, then spread to other parts of the fuel layer. As a result, the problem becomes more complicated: rather than simply determining when the layer is unstable, or following a 1-D burst, a full analysis has to take into account the full 3-D propagation. The numerical problems are extremely tough and it is not yet possible to follow more than about the first 0.1 ms of a burst.

In summary, the burst oscillation phenomenon does seem to demonstrate that burning is nonuniform and propagates, but no one has even attempted a detailed model of this. That's important, too, because many of the specifics of how the burst oscillations occur aren't understood at all, and on their face appear to contradict the simplest ideas. There is, however, still hope that study of these oscillations will have many profound implications.

For example, the light curve we observe depends on the general relativistic deflection of light, and also on the size of the star, so it may be possible (with precise fits) to determine M and R separately for individual neutron stars; indeed, this is the focus of the soon-to-be-launched *NICER* mission, for which I am a team member. You'll recall that precise and accurate M and R measurements will also tell us about the high-density equation of state, so this touches on fundamental nuclear physics. In addition, the propagation of nuclear burning is a phenomenon that occurs in other areas of astrophysics, most notably Type Ia supernovae. Despite current difficulties, nuclear propagation is much easier to study in X-ray bursts (where the matter is kept on the surface) than it is in an explosion such as a supernova. There is long-term hope, therefore, that if we understand nuclear propagation in bursts, we can apply that knowledge to modeling of Type Ia supernovae, and from there have a better idea of the implications of supernovae for the acceleration of the universe and other cool cosmological topics.