Evidence for black holes

Reference webpages:

http://mcdonaldobservatory.org/news/releases/2009/0202.html
and Chapters 8 and 9 in Thorne.

Questions to keep in mind are:

1. What is the role of prior assumptions in acceptance of explanations?
2. How much evidence would be required to convince you of something you disbelieve very strongly?
3. What are common misconceptions about black holes?

Introduction

The simplicity and expected properties of black holes are intellectually exciting, but do they have any relation to reality? Despite all of the mathematical advances in our understanding of black holes through the early 1970s, no one could point to any individual object and say "that has to be a black hole". Indeed, if you think about it, it seems rather difficult: when it comes to black holes, what goes down doesn't come back up. How, then, could we see them?

In this lecture we will discuss the observational evidence that black holes really exist. Our task is difficult: black holes are exciting and exotic, hence we need to point to sources or phenomena that cannot be explained by anything more mundane. Difficult though this is, evidence beginning in the 1960s (or earlier, depending on how you date it) has by now convinced people that most galaxies host supermassive black holes, and we also have good evidence for black holes that are born during the demise of massive stars. We will walk carefully through the evidence; see what you think!

Stellar-mass black holes

In a somewhat ahistorical way, we will begin with the evidence for stellar-mass black holes because it is easier to explain than the evidence for supermassive black holes. To begin, we need to think about how we could see black holes at all. Hawking radiation is negligible, so our only hope is to detect black holes based on their influence on other things. In practice, this means either their gravitational influence on stars (which will turn out to cinch the case for supermassive black holes) or on gas that spirals around them. You might think that black holes could act as gravitational lenses for light sources behind them, and in fact there are proposals for advanced instruments that might see such effects within a few years, but this was well out of reach forty years ago.

If we are considering stellar-mass black holes, of the type thought to evolve from very massive stars, they have to be in a binary for us to see them. The reason is that interstellar gas is far too diffuse, and stellar-mass black holes are too small, for the holes to pick up a significant mass supply that way. Therefore, we need to see a black hole orbiting another star. If we ever get lucky enough to see a pulsar orbiting a black hole then we will be able to get very precise information, but for now we have to deal with black holes in binaries with main sequence stars or giants. To make it more difficult, the distances in astronomy are so enormous that we cannot get a useful image of such binaries; they all look like points to us.

If the companion star is high mass (say, at least several times the mass of our Sun), then it is so bright that the force of the radiation is enough to drive away a significant amount of mass in a wind. If instead the companion is low mass (say, our Sun's mass or less) then its winds are too weak for significant mass transfer. In that case, the only option for mass to flow to the hole is if the companion is close enough that its outer parts flow over to the hole (this is called Roche lobe overflow). The matter has the angular momentum of the orbit, so it does not fall directly onto the hole. Instead, it spirals gradually towards the hole in an *accretion disk*. The closer it gets, the more energy it releases, and near the hole the gas gets up to temperatures of more than ten million degrees. Gas this hot emits mainly in X-rays, which are absorbed by our atmosphere (fortunately for us!). It therefore took until the development of sounding rockets (in the early 1960s) and even more so the launches of X-ray satellites in the 1970s for these sources to be observed. The two categories of binaries are called high-mass X-ray binaries (or HMXBs) and low-mass X-ray binaries (LMXBs).

Easy enough, right? Unfortunately it is more complicated than it may seem, because all the X-rays from the disk tell you is that the gas is very deep in a gravitational well. As we discussed last time, the gas will only spiral down to the innermost stable circular orbit before it drops quickly down. In this respect, the disk is very similar whether the central object is a black hole or a neutron star (as long as the neutron star doesn't have too strong a magnetic field). You might think that there would be a big difference between the plunge to nothingness into a black hole and the hard landing on a neutron star, but in fact the differences are really minor. There are ways of telling that an object is *not* a black hole, but (to gloss over a long and complicated story) there is no signature that has been found in all accreting stellar-mass black holes, and only in stellar-mass black holes.

We therefore have to look for another way to rule out neutron stars if we wish to be certain of having a black hole. To do this, we recall that neutron stars have a maximum mass, just as white dwarfs do. The exact value of this maximum mass is not well known, but very general arguments tell us that it cannot be more than three times the mass of the Sun. Therefore, if we see a bright X-ray emitting object with more than three times the mass of our Sun, it has to be a black hole.

That leaves open the question of how we weigh something that far away. Luckily, we can appeal to Kepler's laws, in particular his third law. In his original version this said that the orbital period in years was equal to the semimajor axis in astronomical units raised to the 3/2 power. For other systems, with different masses, the proportionality is different, but this means that we can figure out the mass if we know the orbital period and the semimajor axis. Unfortunately, we can't measure the semimajor axis directly because we can't get resolved images of these systems. What we can do, though, is measure the component of the orbital velocity that is in our line of sight, via the Doppler effect. More specifically, if we do optical measurements of the companion star and its spectral lines in particular, then the exact wavelength of those lines will shift back and forth during the orbit, giving us both the orbital period and the speed (based on how much the lines shift). It turns out that Roche lobe overflow systems are very nearly circular, so that eases our task.

There are still two ambiguities that remain. One has to do with the way the orbit is oriented. Considering the orbit to lie in a plane, think about the axis that is perpendicular to the plane, and suppose that this axis makes an inclination angle i relative to our line of sight. If $i = 0^\circ$ the orbit is face on, and thus it has zero line of sight velocity even if it is actually moving fast. If instead $i = 90^\circ$, then we see the full orbital speed in our line of sight. A bit later we will talk about ways to estimate i, but in most cases for black hole candidates (as for extrasolar planets!) we don't have a good idea about i.

The second ambiguity comes from the mass of the companion itself. If the companion mass is large, the system will have a high orbital speed even if the compact object is low mass (e.g., if it is a neutron star). If instead the companion mass is small, the compact object mass dominates most of the mass of the system.

With that said, it turns out that if you measure a line of sight velocity of the *companion* of K_1, and an orbital period $P_{\rm orb}$, then we can construct a combination called the *mass function*:

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

Here G is Newton's universal gravitational constant, M_1 is the mass of the companion (which we might or might not know), and M_2 is the mass of the compact object (which we would like to know). Our criterion for the compact object being a black hole is $M_2 > 3\ M_\odot$ (here $M_\odot$ is the symbol for a solar mass).

Looking closely at this expression we can see that f is the *minimum* possible mass for M_2. If $\sin i < 1$ or if $M_1 > 0$, then $M_2 > f$. We also see that f can be computed just from directly observed quantities (K_1 and $P_{\rm orb}$). It is therefore a workhorse in the search for black hole candidates. Finally, we see that if the companion has a high mass (M_1 is large) then f will tend to be small and thus our best option for $f > 3\ M_\odot$ is actually to search for systems that have low-mass companions.

Cygnus X-1 and its successors

Even though LMXBs are theoretically better for establishing black hole candidacy, a high mass binary known as Cygnus X-1 (the first known X-ray source in the constellation of Cygnus, the Swan) was the first and probably still best known stellar-mass black hole source. It was discovered in 1964 during a rocket flight of X-ray instruments, and has been studied extensively since. The companion star to the likely black hole is very bright, and has a probable mass between 20 and 40 times the mass of our Sun. This high mass means that the mass function is quite small, only about 0.25 times the mass of our Sun, but modeling of the companion suggests that the compact object is a few times more massive than a neutron star could be, so it is an excellent black hole candidate. By the way, as of March 14, 2010 the Wikipedia page on Cygnus X-1 makes the claim that "Measuring periodicities in the X-ray emission near the object has yielded a more precise value of 8.7 ± 0.8 solar masses." Don't believe such claims; we don't know enough about the variation of the X-ray emission to say this with any confidence at all!

The low mass function of Cyg X-1 meant that many people, although leaning towards the system containing a black hole, were not absolutely convinced that stellar-mass black holes had been found. This has been resolved over the last couple of decades because quite a few (now more than 20) low-mass X-ray binaries have now been discovered with mass functions greater than the neutron star upper limit of 3 solar masses. The small mass of the companion means that the uncertainty in the black hole mass is reduced. In a few cases, it has even been possible to constrain the inclination i of the orbit. The reason is that stars whose gas flows over to a hole are distorted. Therefore, as the system orbits around and we see the companion from different angles, the amount of light from the companion (as opposed to from the accretion disk) changes in a way that depends on i. We have therefore gone from not knowing any black hole candidates, to having a few where doubt exists, to having many where all reasonable people agree that black holes are by far the best explanation.

It is nonetheless worth pausing to reflect about *why* we should accept the black hole explanation. Black holes are weird things; how much evidence is needed to agree that other explanations don't work? After all, it is always possible to come up with ad hoc objects, and people have done so ("boson stars" being one of the weirdest categories). Fundamentally, this is related to how much evidence is needed for any claim: if the claim is mundane then moderate evidence is fine, but "extraordinary claims require extraordinary proof" to quote Carl Sagan. Black holes are objects that are theoretically allowed to exist. Nothing else that we know to exist (white dwarfs, neutron stars) can explain the data in these systems. Suggested alternatives (boson stars, Q-balls) are not known to exist. Therefore, black holes fit the data and are the best explanation of the phenomena. Nonetheless, a lot of the definitive characteristics of black holes (such as event horizons) have yet to be observed, so there is still a lot of room for greater understanding.

It is also worth stepping back to realize how rare it is that a stellar-mass black hole makes itself visible. We mentioned twenty-some LMXBs that appear to have black holes in them. But given that about 1 in 2,000 stars is thought to become a black hole, and there are 200-400 billion stars in our galaxy, that means that there are probably a good 100 million stellar-mass black holes in our galaxy, of which we are confident of only about 20! Clearly, a very special set of circumstances is needed to make such a black hole apparent to us.

Supermassive black holes and active galactic nuclei

We now go back in time again. For almost all of the history of astronomical observation, we have been limited to what we can learn from visible light (also called optical light, in the wavelength range of 400 nanometers to 700 nanometers, roughly). In the previous section we learned about the information from X-ray observations, but the first big extension (not counting minor observations of near infrared radiation) occurred serendipitously in the 1930s when Karl Jansky observed radio waves from the heavens. After a hiatus due to World War II, radio astronomy got into full swing, with both larger receivers and the clever technique of interferometry. In this technique, separate radio telescopes that are widely separated from each other combine their observations so that their angular resolution is as good as if they were a single telescope with an aperture equal to the separation (but not nearly as large a collecting area, of course). People are working on such things for optical light, but atmospheric fluctuations are much faster at short optical wavelengths than at long radio wavelengths, so it is much more difficult in optical.

People began to catalog individual radio sources, and after a while it became clear that one category of such sources was rather puzzling. These were sources that were smaller than the angular resolution, even as the angular resolution got better and better. Stars are like that, so these were called quasi-stellar radio sources, or quasars. Given that things outside our galaxy (e.g., other galaxies!) are definitely extended objects, people were stuck in a mental block and tried to fit the strange spectra of quasars as if they were stars in our galaxy (leading to discussions of uranium lines and the like!). Eventually, in 1963 the young astronomer Maarten Schmidt at Caltech realize that one particular source, 3C 273, was actually a distant object: its redshift was the largest then known, and it had an implied distance of more than two billion light years! Once this discovery was announced, many other quasars were shown to have distances that were even greater. The current record holder sent its light to us when the universe was less than a billion years old.

But now the puzzle was even greater. Quasars are bright electromagnetic sources, not just in the radio but across the whole band. They are also really distant, so to appear as bright as they do they must emit an enormous amount of energy per second. In fact, some quasars shine as brightly as a thousand normal galaxies! To make it even odder, better and better angular resolution failed to resolve them, and the characteristics of their variability mean that the powerhouses of quasars must be confined within a region smaller than our Solar System. How could this be?

After the 1963 discovery that quasars are extremely energetic, a variety of ideas were suggested including that they were powered by supermassive stars (more than a hundred million solar masses). These are unstable, however, and it took until the end of the 1960s for a few people (e.g., Donald Lynden-Bell) to suggest what we now consider the most promising idea: quasars, and other types of energetic "active galactic nuclei" are powered by black holes.

At first glance, this seems a rather odd idea. We think of black holes as surly loners, sucking down matter and giving nothing in return. However, as we indicated earlier, matter does not just plunge straight into black holes (it usually can't; things with angular momentum have to spiral around). Therefore, the gas has friction with itself, heats up, and radiates...a *lot*. For comparison, suppose we took some mass of hydrogen (e.g., a kilogram) and caused all of it to fuse into helium in a hydrogen bomb. This would release 0.007 (James Bond!) of its mass-energy: $E = 0.007mc^2$. By comparison, if that hydrogen were to spiral into a moderately-spinning black hole, it would release $E = 0.1mc^2$ of energy, more than 10 times as much. This is a stable release of energy, and even a black hole that is a billion times the mass of our Sun would fit within the orbit of Uranus, so this satisfies all observational criteria.

These insights are not won easily, and such an exciting conclusion has to be put to rigorous tests. By the mid to late 1990s almost all astronomers were convinced that active galactic nuclei were powered by supermassive black holes. Indeed, a "unified model" had emerged, in which spinning supermassive black holes dragged reference frames around them, producing twisted magnetic fields, and created high speed "jets" (think of a firehose of matter and energy) along the spin axis. Depending on how we look at such a system, we see different types of active galactic nucleus phenomena.

But most galaxies are not active in this sense; in fact, only about 1% are (depending on your definition). It might be that only 1% of galaxies have supermassive black holes, it might be that all galaxies do but are active just 1% of the time (maybe they have an intermittent fuel source), or maybe it is something in between. How could we tell?

In the mid 1990s, groups led by Andrea Ghez at UCLA and Reinhard Genzel at the Max Planck Institute for Extraterrestrial Physics (Garching, Germany) started using infrared observations of stars very near our galactic center to determine if those stars were orbiting around a very massive but unseen object. In a way similar to how we get masses for stellar-mass black holes, their observations over the last 15–20 years have demonstrated that we *do*, in fact, have a supermassive black hole. Not surprisingly, given that we are very close to our own galactic center, the case for a black hole in the Milky Way is essentially beyond doubt, whereas more distant galaxies have poorer data (but are still convincing to anyone who wasn't bitten by a black hole as a child). Ours has about four million times the mass of the Sun; the biggest known have about three *billion* times the mass of our Sun. Pretty impressive!

Black hole – galaxy coevolution

Also in about the last 15–20 years, people have realized that the properties of supermassive black holes are correlated with the properties of the galaxies around them. For example, bigger galaxies have bigger black holes. This is not all that surprising in the abstract, but what is remarkable is that the correlation between the masses is so strong. In fact, it is such a good correlation that people have concluded that galaxies and their supermassive black holes have to evolve together, and it appears that the activity of central black holes has a major role in how stars form. It probably even dictates important aspects of how clusters of thousands of galaxies evolve; pretty impressive, when you realize that these clusters have a million times the mass of the central black hole!

Black holes have come a long way. Fifty years ago, the term was not used and most physicists would have been, at best, agnostic about their existence. Now, it appears that black holes are not only common but that they play important roles in governing the most massive structures in the universe. Future observations will be crucial to determine whether black holes really obey the laws of general relativity, and from the astrophysical standpoint we need to learn a great deal more about how supermassive black holes originate and evolve. Both topics will be strongly aided by gravitational wave observations, which will be the topic of the next lecture.