Crashing Black Holes

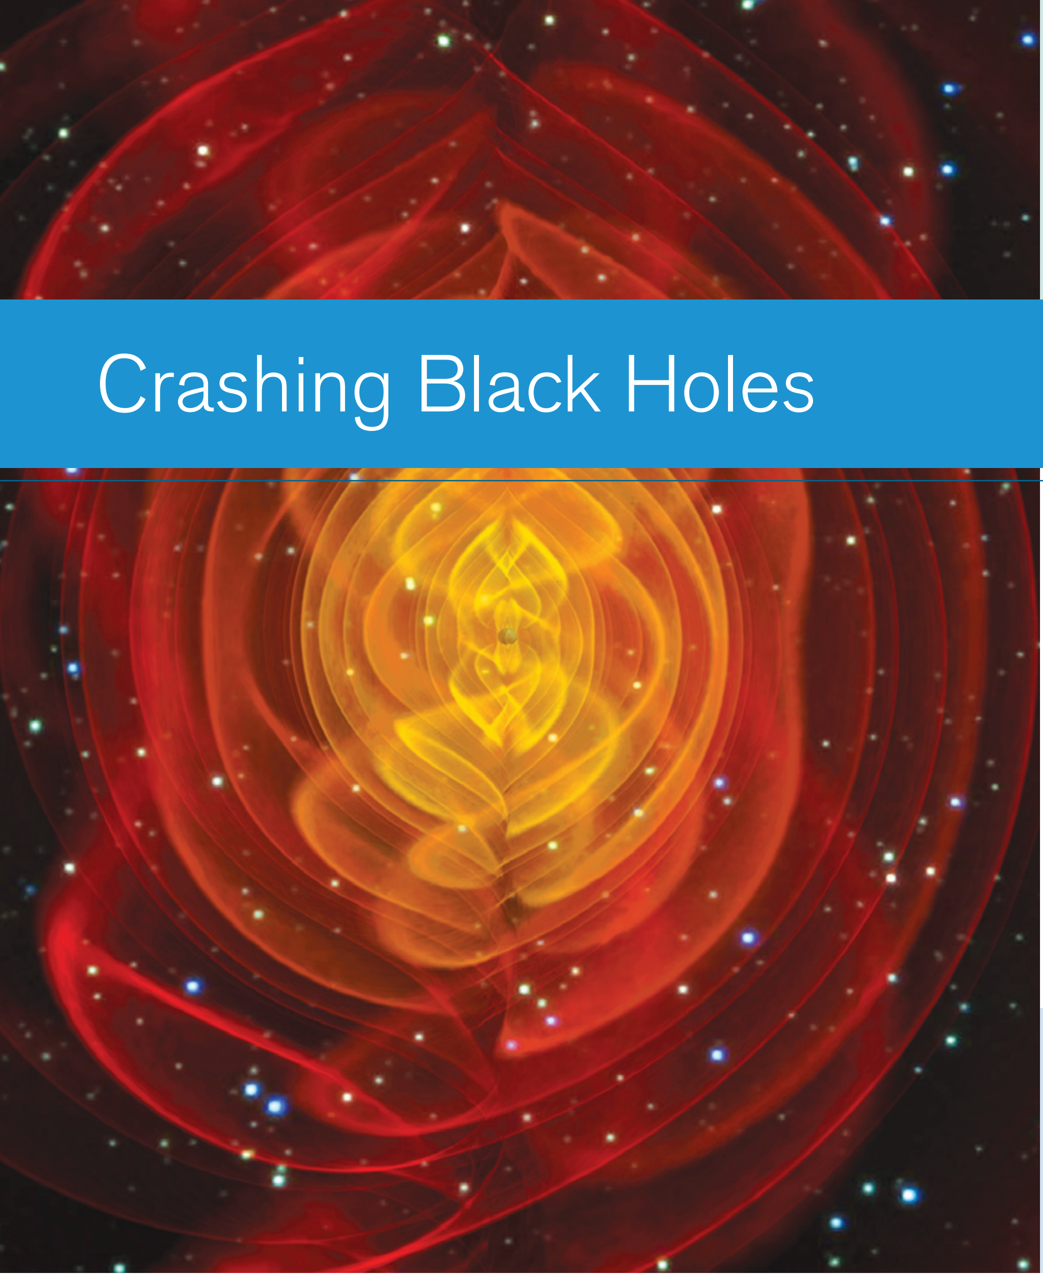

A still from a movie of two black holes colliding. Taken from a simulation by the team at NASA Goddard, this image shows the merged black hole in the center and what the resulting gravitational waves would look like (in orange).

After struggling for years, physicists have now succeeded in simulating black-hole collisions, heralding a new era in physics.

On Thanksgiving morning in 2004, Frans Pretorius made a few last-minute changes to his computer code. Pretorius, a Caltech postdoc, was struggling with a problem that had baffled theoretical physicists for decades—simulating the collision of two black holes. A black hole, whose gravity is so strong that even light cannot escape its pull, is indeed a strange creature—but not a rare one. Black holes are the inevitable fate of massive stars, and since those stars often exist in paired systems called binaries, astronomers suspect that black-hole binaries are scattered all over the cosmos. As two black holes whirl around each other at nearly a third of the speed of light, they create ripples in space-time, the fabric of the universe. To Earth-bound physicists, these as yet undetected ripples, called gravitational waves, will tell the story about the black-hole pair's dance and ultimate union, lending insight into how gravity—and the universe—works. But to read the ripples, you'll need simulations to know what they'll look like. Unfortunately, the math behind the science is so complicated that no computer had ever run for more than a fraction of an orbit before the inevitable program crash. Pretorius, however, had an inkling that this time might be different.

He was applying an approach developed a few years prior by David Garfinkle, now at Oakland University in Michigan. By Thanksgiving, Pretorius's code was still not working. But the night before, Carsten Gundlach, a visiting researcher from the University of

By Marcus Y. Woo

The idea of a **black hole** actually goes back to 1783. John Michell, a British natural philosopher and geologist, reasoned that a "dark star" could exist if light couldn't escape its gravity. Newton's laws showed that a star's escape speed—the speed needed to leave its gravitational pull—is proportional to the square root of its mass divided by its radius. So if the radius were sufficiently small for a given mass, the escape speed could be greater than the speed of light.

Southampton in England, had proposed a few more minor tweaks. "I decided that if I had a few minutes free, I'd try it," Pretorius recalls. Because the problem is so complex, he ran the revised program with the simplest possible test case—a single, stationary black hole—and even this scenario was far from trivial.

He tapped into the power of the Lonestar cluster at the University of Texas, Austin, which combines 5,200 processors to form a supercomputer. Even with such muscle, the simulation would take hours to run, giving him and Gundlach ample time for a holiday dinner at a fellow researcher's house. On Sunday morning, when the simulation was done, he scanned the numbers. "At that moment, I knew it would work," he recalls. "That was very exciting." Unlike all the simulations tried before, his was stable. And if it worked for one black hole, it should work for two.

Over the next six months, Pretorius put a binary system through one full orbit with a merger at the end. He then published his results in *Physical Review Letters*. Six months after that, two other groups, from NASA Goddard Space Flight Center and the University of Texas, Brownsville, independently published their own work, heralding a new era in gravitational physics. Pretorius, however, was not only the first to succeed, but his method was unique—and he did it largely on his own.

"It was really an outstanding piece of work," says Lee Lindblom (BS '72), a senior research associate in theoretical astrophysics. The intellectual genealogy of the Goddard and Brownsville groups originated at the Albert Einstein Institute in Germany, Lindblom says, so their methods were similar. "Frans came from left field—he didn't come from that group—and used a rather different approach, a different version of Einstein's equations. And as soon as the other groups realized this could be done by this upstart guy, I think they went crazy and worked 24 hours a day. I think they were pretty close, and in a few months, they were able to put it all together."

Since then, groups from around the world have gotten in on the act. Pretorius has moved to Princeton and into different fields, but the Caltech team, led by Lindblom and Mark Scheel, a senior research fellow in physics, and under the oversight of Feynman Professor of Theoretical Physics Kip Thorne (BS '62), and a Cornell group, led by Thorne's old student, Saul Teukolsky (PhD '74)—together, the largest numerical relativity group in the country—carry on. Their simulations are the most accurate of any so far, orders of magnitude beyond Pretorius's initial work. Ultimately, they'll be able to recreate merging black holes of any size, spin, and orbit. But this work wouldn't be where it is had it not been for one late night in November 1976.

That evening, Thorne found solace in the streets of Pasadena. He needed fresh air and he needed to think, debating with himself whether to propose that Caltech get involved in an ambitious new gravitational-wave detector. The risks, as he later noted in *Black Holes and Time Warps*, were enormous. There was no guarantee of success, and the project would be expensive—not only financially but also in terms of energy and time. But, after months of internal struggle, he went for it. Almost two decades later, LIGO, the Laser Interferometer Gravitational-Wave Observatory, spearheaded by Thorne; Ronald Drever, now professor of physics, emeritus; and MIT's Rainer Weiss, was born.

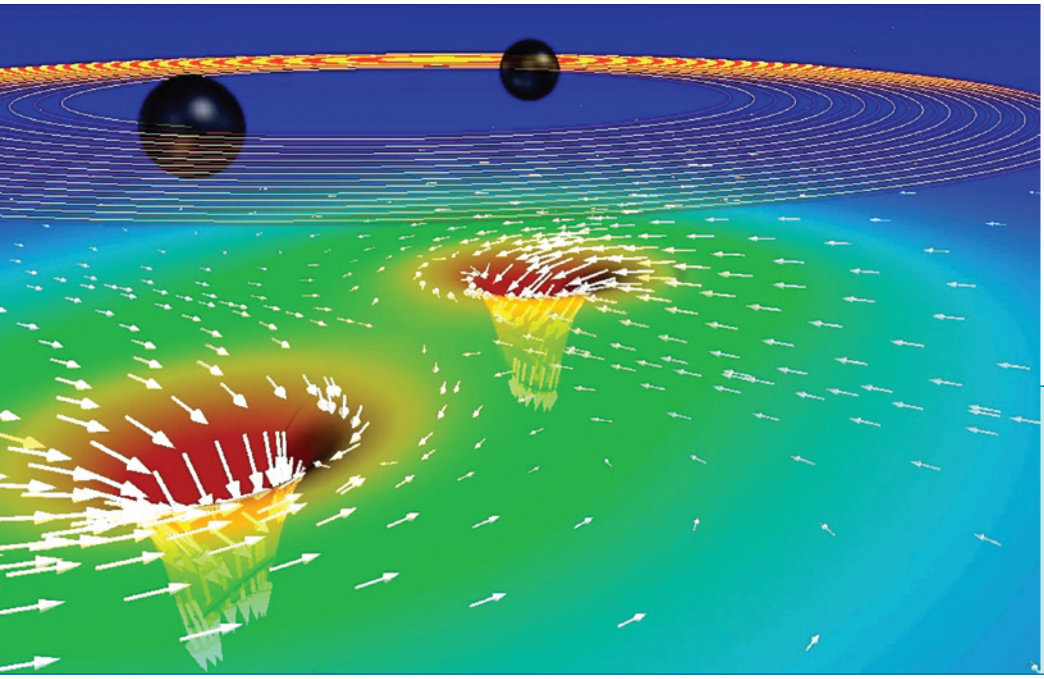

This snapshot from one of the Caltech–Cornell team's simulations depicts two black holes prior to merger. The colored surface represents the curvature of space-time. The arrows represent the flow of space, similar to the flow of water in a whirlpool.

An aerial view of the LIGO facility in Livingston, Louisiana.

ERGO, LIGO

In 1916, Einstein declared that mass or energy bends space and warps time, and this curvature of space-time is what we experience as gravity. Newton's apple falls not because the ground is pulling it down, but because Earth is curving the space-time around it. The apple is just following the bends of space-time.

Any bit of moving matter—say, coalescing black holes, an orbiting Earth, or even a gesturing hand—creates waves of space-time that propagate across the cosmos just as a tossed rock sends ripples across a pond. Suppose you arrived at the pond an instant later. In principle, you could look at the ripples' patterns and deduce the rock's size and trajectory. A softball-sized one would splash while a pebble would barely disturb the shimmering surface. In similar fashion, gravitational waves betray their source. And it turns out that the most likely source, colliding black holes, may be fairly common.

> The Milky Way is whizzing toward the Andromeda Galaxy at 120 kilometers per second, and the two are slated to meet in a few billion years.

Although the precise number is hard to pinpoint, a large portion of stars exist in binary systems. In fact, astronomers like to joke that three out of every two stars is a binary. Many of the universe's binaries are pairs of massive stars, and if they're at least about ten times heavier than the sun, they'll collapse into black holes once they've exhausted their fuel. As the brand-new black holes continue to revolve around each other, gravitational waves carry away energy and angular momentum, sending the holes into a death spiral until they crash and trigger an explosion of gravitational waves.

While these signals are the target for LIGO, there's another kind of gravitational-wave factory that makes space-time ripples of lower frequencies, ripples likely to be visible only to a future space-based observatory called the Laser Interferometer Space Antenna (LISA). These factories are merging supermassive black holes, behemoths millions to billions of times more massive than the sun, and they would set off the strongest gravitational waves.

These humongous holes inhabit the centers of galaxies, which fly through the universe like Frisbees. Once in a while, galaxies get too close, succumb to their mutual gravity, and smash into each other. Amid the swirling stars and gas, the two giant holes find each other and merge. In fact, our own galaxy, the Milky Way—which has its own black hole of three million solar masses—is whizzing toward the Andromeda Galaxy at 120 kilometers per second, and the two are slated to meet in a few billion years. At Andromeda's center awaits a black hole almost 50 times bigger than ours.

Still, there's no way to know for sure how many black holes there are, since by definition you can't see them. "The uncertainty is one of the reasons why the first detections with LIGO will be so interesting," says Stan Whitcomb, chief scientist for LIGO. He says that given what's known about stellar evolution and LIGO's sensitivity, the instrument should detect anywhere from roughly one merger per century to one per year. Although that may not sound too promising, upgrades planned for 2014 will boost those rates more than a thousandfold.

Like pond ripples, gravitational waves

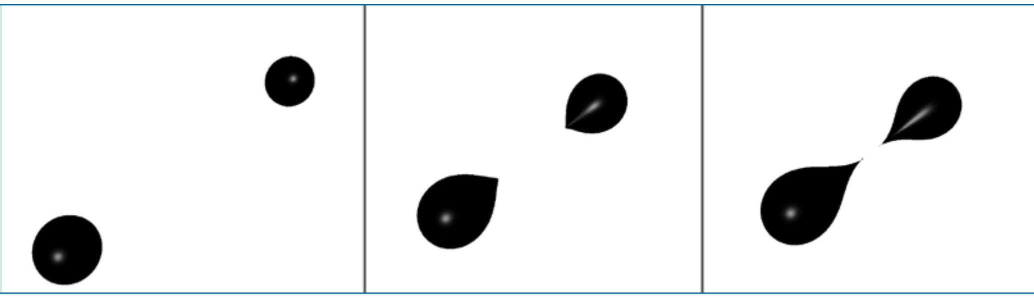

These time-lapse snapshots show two black holes on a direct collision course. The black holes' surfaces are called the event horizons. Note how the horizons extend toward each other before merging.

In general, Einstein's equations are impossible to solve without a computer. But there are a few simple cases in which all you need is pencil and paper. In 1915, the same year Einstein introduced his theory of general relativity, an astrophysicist named **Karl Schwarzschild** solved the equations for a stationary sphere—a nonspinning black hole. Not only did he do it by hand, he did it while fighting for the German army on the Russian front, where he would die from an illness the following year. His work, however, lives on as an important contribution to science.

weaken over distance, and they may travel millions of light-years to get to Earth. When they do, even the strongest waves will stretch and squeeze space-time by no more than one part in 10^{21}. Measuring such puny waves, then, takes breakthrough technology.

LIGO has to detect motions as small as 10^{-16} centimeters—a thousand times smaller than a proton. As its name suggests, the observatory is an interferometer, an L-shaped structure whose arms stretch four kilometers. A laser at the junction fires half its beam down each arm. To keep stray gas molecules out of the beams, the beams are encased in steel pipes 1.2 meters in diameter, evacuated to one-trillionth of normal atmospheric pressure. At about 8,000 cubic meters, the vacuum system is one of the biggest in the world. At each end of each arm is a mirrored test mass, a 10-kilogram cylinder of ultrapure fused silica. The two beams bounce back and forth between the test masses, and are eventually recombined where the arms meet. If all's quiet in the universe, the light waves will match up peak to peak and trough to trough. But if a cosmic cataclysm sends a gravitational wave our way, the arms will alternately stretch and shorten, shifting the laser beams out of phase with each other. Because the shift is minuscule, the beams bounce back and forth about 100 times to accumulate a difference big enough to measure.

After nearly a decade of construction, LIGO went online in 2002. There are two nearly identical observatories—one in Hanford, Washington, and one in Livingston, Louisiana. Having two detectors allows you to spot where the waves are coming from. Double detections also help confirm a signal; if both say they saw a wave, then chances are they really did.

You need to screen out false detections because LIGO is extremely sensitive, capable of picking up ocean waves crashing onto a rocky shore hundreds of miles away, and even the stealthy shifts of creeping geological faults. According to Lindblom, researchers even detected an airplane flying overhead. They saw the vibrations caused by the sound of the plane's engines and measured its Doppler shifts—the shifts in frequency that depend on the plane's motion relative to an observer—and compared them with known flight plans to confirm that they had indeed detected the airplane.

As in the airplane example, the key to discerning sources from noise—and from each other, such as distinguishing between systems with different masses—is to know what you're looking for. Picking out a merging black-hole binary is akin to picking out a single voice in a noisy restaurant. Just as cupping your ears won't necessarily help you eavesdrop on dinner conversations, LIGO's great sensitivity won't necessarily help if you can't interpret what you see. And that's where numerical simulations come in.

NUMBER TRICKS, NINJA, AND KICKS

The mid-1990s saw vigorous efforts around the world to simulate black-hole mergers. But as the 20th century came to a close and LIGO came online, progress was slow. "At that time, nothing worked," Lindblom remarks. "The field was a mess." Success was measured in terms of fractions of an orbit, Pretorius says. Physicists could get the black holes to begin their dance, but if the two spiraled in too close, the code crashed before they could consummate their union. Each fix would extend the simulation a bit farther along, but the crash was inevitable—solving one problem just revealed another one. Taxpayers had invested a lot of money in LIGO, and the space-based version, LISA, was already in the works. Thorne became worried.

"I decided that we really needed to become involved to try to push things forward," Thorne says. So, in 2001, Caltech plunged in with the group at Cornell, which had been working on the problem since the 1980s. He began building a team, and he recruited Lindblom, who had been splitting his time between Caltech and Montana State University studying neutron stars (neutron stars are black holes' less-dense cousins), and Mark Scheel from Cornell as its leaders. The team attacked the problem head-on.

The challenge isn't the physics, but the math and how to program it. Some of Einstein's equations are called constraint equations, which are conditions that the solution must satisfy at all times. Consider a marble rolling on a table. Newton's equations of motion—the ones you learned in high school—describe how the marble moves around. The constraints confine the marble to the table's surface.

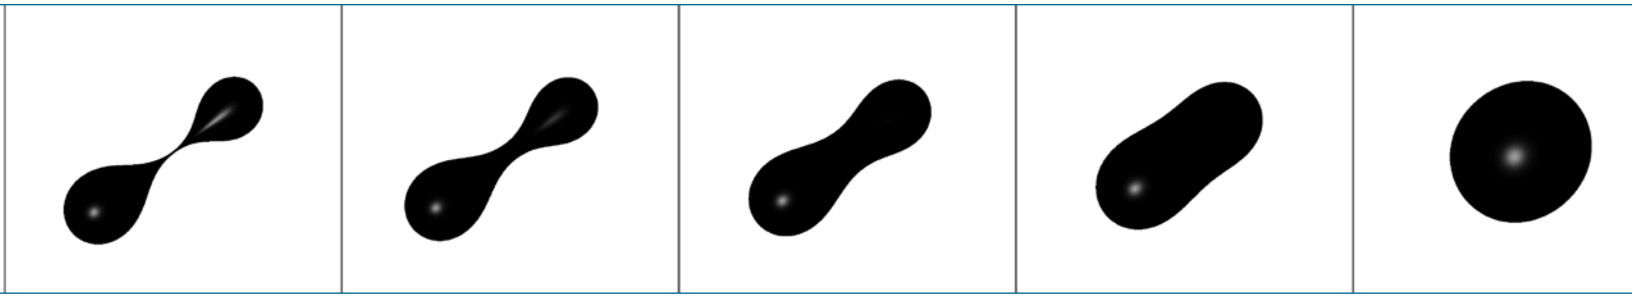

“In the last few years, it’s been about the physics. It’s very exciting.”

Solving Einstein’s equations, of course, isn’t exactly easy. There are 10 equations with usually thousands of terms, and just keeping track of all the numbers is practically impossible. As powerful as computers are, they’re far from perfect. For example, numbers such as π or 1/3, whose decimal forms go on forever, must be truncated to fit in the memory. The discrepancies between the true and truncated values pile up, causing the errors to grow exponentially until the constraint equations are violated, and the marble flies off the table.

“People realized that the standard way of writing Einstein’s equations was bad,” Lindblom says. “That’s one of the things where our group made a lot of contributions, in figuring out what was killing everybody. Then Frans [Pretorius] and the other groups found stable ways of writing the equations, and once they did that, they were up and running.”

The trick was a mathematical method called constraint damping. Consider a constraint equation that requires a certain term to be zero. You can add that term to the equations and not change anything, because you’re adding nothing. Then you can rewrite the equations such that when errors cause the constraint term’s value to change, the new terms push the constraint back toward zero. In other words, the program has a negative-feedback system to keep the violations down. Figuring how to do that, Thorne says, was a mathematical tour de force.

Another obstacle is the singularity that lurks in the black hole’s belly. A singularity is a point where space-time has collapsed to infinite gravity—and infinite curvature. The mathematical equivalent of dividing by zero, it’s where Einstein’s theory of general relativity breaks down. But the researchers realized that what happens inside the black hole doesn’t matter.

A black hole’s voracious appetite only comes into play when you get really close to one. From a distance, they behave as any star or planet would. If the sun were to be replaced by a black hole of the same mass, Earth would not be any different—aside from the lack of sunlight and resulting mass extinctions. As you get closer to a black hole, however, the beast begins to reveal itself. Once you cross the so-called event horizon, you’ve been swallowed—the gravitational pull is so strong that not even light can escape. This, then, is the black hole’s “surface,” beyond which communication with the outside universe is impossible. What happens inside the black hole stays in the black hole, so the program need only run until the event horizon is encountered, avoiding the messiness of the singularity.

The black holes move in their orbits, of course, adding another challenge to the simulations. The computer code imposes a grid onto the black hole and calculates how much space-time is curved at each of the grid’s points. Other groups had kept the grid stationary while the hole moves. But the Caltech–Cornell team was devising a more accurate method, which needed the grid to follow the hole’s motion, Lindblom explains, so the researchers had to invent techniques to glue the coordinates to the black holes.

Which brings us back to Pretorius’s Thanksgiving weekend in 2004. His achievement triggered an exceptional period in the field, Lindblom says. Simulations featuring multiple orbits are now routine, and after years of doing nothing but programming, scientists are returning to the physics that originally drove their research. “It was very disheartening that we spent so much time on just coding and mathematical analysis,” Pretorius says. “In the last few years, it’s been about the physics. It’s very exciting.”

Armed with the new methods, researchers, including those at Caltech, have been involved in a game of hide-and-seek witih each other involving virtual binary black holes. This experiment, called NINJA, for Numericial INJection Analysis, is designed to find out how good physicists are at spotting black-hole pairs amid LIGO’s background noise. One team submits a simulated data set, and the data are then buried in statistical noise. Other teams then sift through the noisy data to try to determine the masses, spins, and orbital shapes of the black-hole binaries. So far, the seekers are winning.

On another front, researchers at the Rochester Institute of Technology have shown that colliding supermassive black holes can trigger a “kick” that sends the coalesced monster careening through space at up to 4,000 kilometers per second. Furthermore, they’ve found that the speed and direction of the kick depends on the

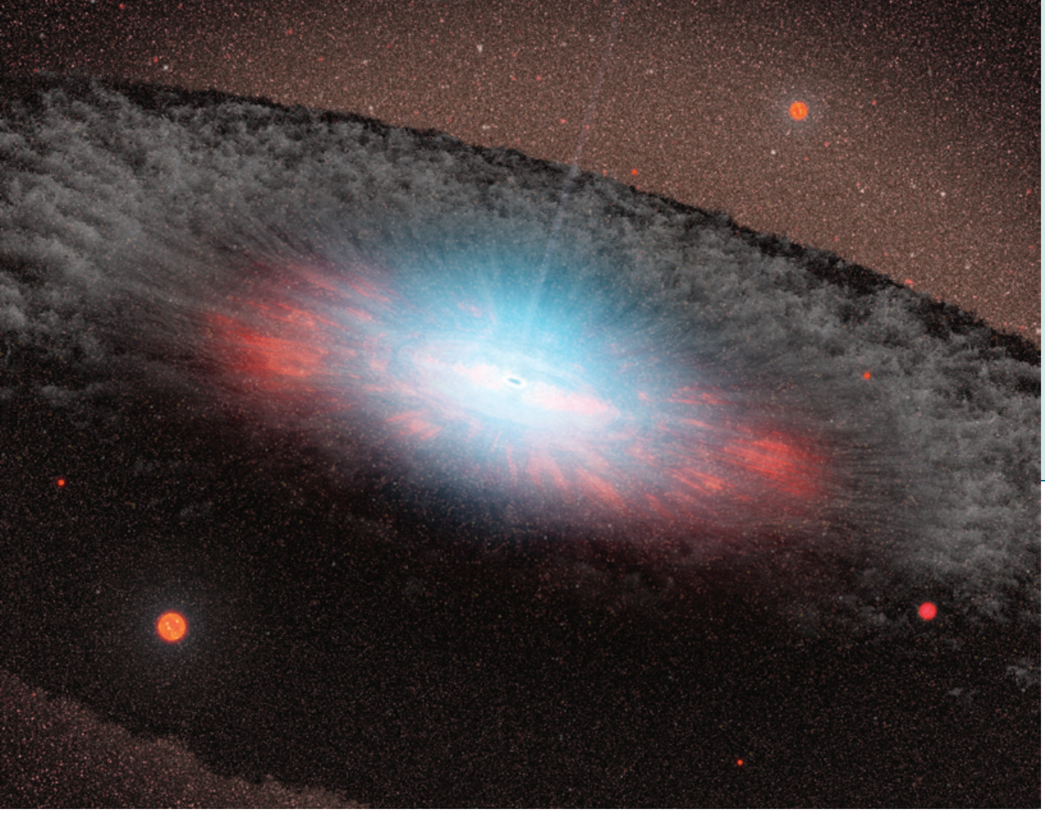

An artist's conception of a supermassive black hole at the center of a galaxy. The gray matter is a torus of gas and dust, and as it's devoured by the black hole, tremendous amounts of heat and radiation are released, represented by the blue color.

pair's initial spin alignment—and not their masses. This surprising result could help pin down the rules of how supermassive black holes form. Astronomers believe that nearly every galaxy harbors one at its center, and some scientists have suggested that the hole is formed through multiple mergers. But 4,000 kilometers per second is more than enough to escape a galaxy's gravity, so mergers may not be an efficient way to build these gravitational goliaths.

SIMULATING EXTREME SPACE-TIMES

The Caltech–Cornell team, however, wants to study the gravitational waves from merging black holes in unprecedented detail, with the high accuracy that LIGO needs. To do this, the team decided to pursue a numerical technique that is the most accurate yet—but also more unforgiving.

All previous calculations had relied on so-called finite-difference methods. Consider again the marble on the table. The equations that govern the marble's motion describe its changes in position over time. Mathematically, this change is infinitesimally small from moment to moment. But a computer doesn't know what "infinitesimally small" means—it needs an actual number. Finite-difference methods approximate these changes as tiny, discrete steps, so that the marble would roll around in steps of, say, 0.1 millimeters.

But there's another technique, called a spectral method, that exploits the smoothness of space-time. The curvature around black holes changes gradually—there are no sudden jumps or shocks. Roughly speaking, spectral methods approximate the answer as the sum of a series of well-understood mathematical functions. The more functions in the series, the more accurate the answer. With this technique, the calculation converges toward the exact answer far faster than with finite-difference methods.

Of course, there's a catch. Spectral methods are more complicated and delicate—so much so, in fact, that the SXS project (for Simulating eXtreme Space-time), as the Caltech–Cornell effort is known, and two smaller groups in France and Germany are the only ones in the world who use them. Finite-difference methods allow for some leeway, Lindblom says. Even if the mathematics isn't strictly correct, finite-difference codes will still give you an answer that's not terribly wrong. "But spectral-method codes just tend to blow up. They just won't compute anything unless you've done every single thing right."

Only in the last year have they succeeded, Thorne says. "We are now doing mergers and getting very high accuracy waveforms for use in LIGO data analysis—waveforms that have higher accuracy than LIGO is ever likely to need, even decades into the future." But this level will be needed for LISA, which will make more accurate measurements.

Because of the computer code's increased complexity, SXS has yet to fully automate its merger simulations. The researchers are having to stop the program numerous times and adjust it before starting it up again. But, they say, automated simulations of complete mergers will soon be routine. Thorne says, "We're playing catch-up in the sense that [other groups are] able to do a given problem sooner than we are, but they can't get the precisions that LIGO and LISA require."

Because of all the money and time put into gravitational-wave detectors, simulating black hole collisions is the main task at hand. But the physicists haven't forgotten that black holes—with their ability to slow clocks, bend light, and drag space-time around themselves like whirlpools—are the universe's funhouses. Black holes are where weird stuff happens, and the team is using

Discovered by Charles Messier in 1773, the Whirlpool Galaxy lies 23 million light-years away. That's a lot of frequent-flier miles.

their computing techniques to learn all about the bizarre drama of extreme gravity. For example, Caltech's Cornell collaborators are simulating a black hole devouring a neutron star, a gruesome ordeal in which the latter is torn apart and swallowed.

Another puzzle relates to overcoming the obstacle of overly complicated math. As two black holes coalesce, Einstein's equations have to be used in all their gory detail. But when the black holes are still on their approach and extreme gravity has yet to kick in, physicists use the post-Newtonian approximation—a mathematical tool that mixes a bit of Newton's physics with a bit of Einstein's. A bridge between the classical and relativistic depictions of gravity, it avoids the complications of Einstein's equations. The question is when and how the approximation stops describing physics and starts spitting out nonsense. Now, with their high-accuracy simulations, Thorne says, they can find this transition and see how the approximation breaks down. "We can see how far you can trust it and when you have to abandon it."

Part of what makes the region around black holes so weird and complicated is that it's nonlinear—curvature begets more curvature. Thorne likens it to an avalanche, in which falling snow grabs more snow, which in turn accumulates even more snow. Along with Assistant Professor of Physics Yanbei Chen (PhD '03) and a few graduate students, and armed with the Caltech–Cornell team's simulations of spinning, colliding black holes, Thorne is trying to understand how this avalanche behaves. "That's a process we're just beginning, and we see it as a major direction. I don't know any other groups that are pursuing that at the moment, but it's a big thing for us."

Next, Thorne says, they might confront a topic more commonly associated with science fiction—wormholes. A wormhole is like a black hole, in that its powerful gravity keeps photons on a short leash. But whereas a black hole is a dead end, a wormhole is a tunnel that connects two places in the universe, conceivably serving as an intergalactic shortcut for space travelers. But before you plan your trip to the Whirlpool Galaxy, note that some theorems say wormholes would collapse into a singularity, according to Thorne. But no one knows exactly how a wormhole collapses—or if something can prevent the collapse.

LIGO has yet to detect anything, but as in all science, a nondetection is still a result. For example, on February 1, 2007, a gamma-ray burst lit up an area of the sky occupied by the Andromeda Galaxy. Conventional observations had led many astronomers to believe that this particular burst was from a relatively weak source in Andromeda, perhaps a collision between two neutron stars. If such a collision had happened as close as Andromeda, it would produce lots of detectable gravitational waves. But it would also have produced an even brighter burst in electromagnetic waves, which astronomers didn't see. So when LIGO didn't see anything, a neutron-star collision in Andromeda was pretty much ruled out. Instead, the source was probably either a collision farther away, or a so-called soft-gamma repeater—a single neutron star whose strong magnetic field produces rapid, successive bursts of gamma rays, but hardly any gravitational waves. So far, the data suggests the latter is more likely.

Last June, LIGO scientists published the results of a year-long study on the Crab Pulsar, a spinning neutron star whose birth was recorded by Chinese astronomers in 1054. A neutron star is so dense that its atoms have collapsed, its protons and electrons squishing together to make neutrons. Although it's almost a perfectly smooth sphere, a neutron star, with a typical radius of 10 kilometers, can be deformed by a few centimeters. And these tiny bumps are enough to generate gravitational waves as the star whirls around its axis at breakneck speeds. The Crab Pulsar, for example, spins some 30 times per second. The waves carry energy away and slow down the pulsar's spin. Astronomers had already known this pulsar was slowing down, and had carefully measured how fast it was losing energy. But when LIGO didn't detect any gravitational waves, the researchers could set a limit as to how spherical the pulsar could be, concluding that if the pulsar were 10 kilometers in radius, any deformities couldn't be more than one meter in size, beyond what's physically possible. The scientists then showed that gravitational waves could account for no more than 4 percent of its energy loss. The rest must be lost elsewhere, almost certainly through the rapidly spinning magnetic field, researchers say.

Despite these useful results, LIGO will only be truly successful if it sees something. Earlier this year, the National Science Foundation began funding for Advanced LIGO, a seven-year, $200 million project to improve LIGO's sensitivity tenfold—a leap that would mean physicists should see anywhere from 16 to 2,700 mergers per year, according to scientists. And if LIGO still doesn't find anything, we may have to revamp our understanding of gravity—as big a potential shakeup as any in science.

The 1960s and early 1970s have been dubbed the Golden Age of general relativity. During those 15-odd years, physicists at Caltech and elsewhere transformed an intellectual curiosity into a field that continues to inspire. Now, more than three decades later, we may stand at the brink of yet another Golden Age. For physicists, it's about time. e&s

Combining optical and X-ray data from the Hubble and Chandra space telescopes, this image reveals the Crab Pulsar in detail. Shock waves carry wisps of matter, seen in the picture as rings, out at half the speed of light. Jets of radiation shoot out perpendicular to the rings.