Pulsar timing arrays are poised to reveal gravitational waves

Radio observatories are accumulating data to detect mergers of supermassive black holes.

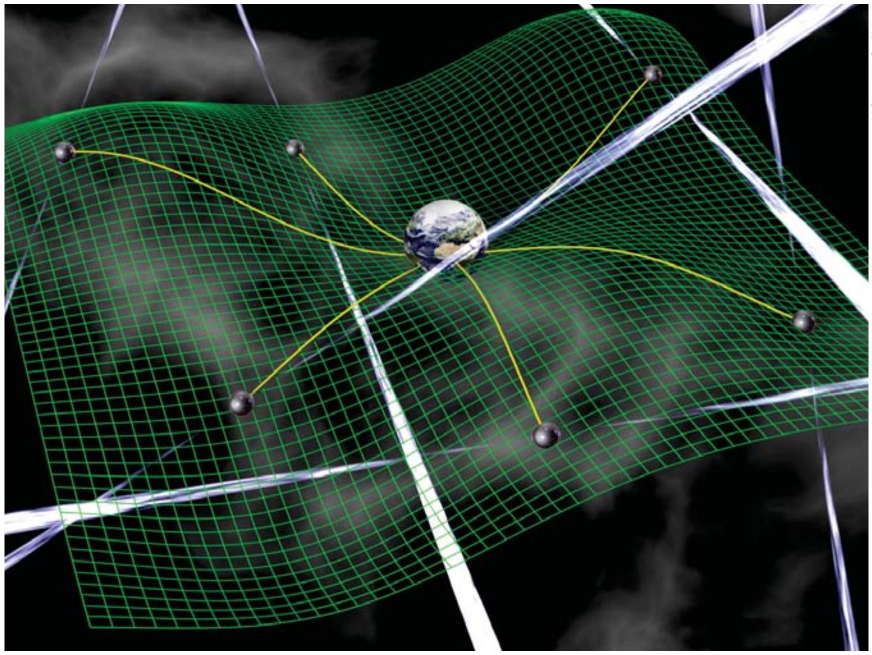

DAVID CHAMPION

GRAVITATIONAL WAVES from accelerating masses distort spacetime (green grid). The distortions squeeze or stretch the travel time to Earth of light from pulsars (black spheres). Scientists are exploiting the regularity of pulsar signals to hunt for gravitational waves.

Pulsars emit radio waves that, unless altered, arrive at Earth like clockwork, and astrophysicists are exploiting that precise periodicity to create a galactic apparatus for detecting gravitational waves. The trick is to tease out perturbations in time-of-arrival data from scores of millisecond pulsars that have all encountered the same gravitational waves.

The pulsar timing array method is most sensitive to low-frequency gravitational waves, and so requires years of observations. "We are now on the edge of when we expect to make a detection," says Xavier Siemens of the University of Wisconsin–Milwaukee.

The first millisecond radio pulsar was discovered in 1982. In the mid 1990s, as more and more pulsars were discovered and classified, their potential as tracers of gravitational waves was recognized. "That's why the different projects all started at roughly the same time," says Joris Verbiest of Bielefeld University in Germany. The three ongoing projects began in the mid 2000s: the North American Nanohertz Observatory for Gravitational Waves (NANOGrav), the European Pulsar Timing Array (EPTA), and the Parkes Pulsar Timing Array (PPTA) in Australia. They cooperate through the International Pulsar Timing Array (IPTA).

A natural apparatus

In February 2016 the Laser Interferometer Gravitational-Wave Observatory (LIGO) and Virgo collaborations announced to great fanfare the first-ever direct detection of gravitational waves (see PHYSICS TODAY, April 2016, page 14); two more observations have since been reported. Ground- and space-based interferometers must achieve extremely precise alignment so that length changes of less than a nanometer can be detected. With the pulsar timing arrays, by contrast, the experiment is "kindly and beautifully provided by nature," Siemens says. The challenge lies in identifying and interpreting teensy changes—a few tens of nanoseconds—in the times of arrival of pulsar signals.

The two approaches are sensitive in different regimes. LIGO detected the death throes of stellar-mass black holes as they spiraled toward each other. The pulsar timing arrays can detect vastly more massive objects that are moving slowly. They are looking for events involving objects on the order of a billion solar masses, such as the supermassive black holes that inhabit the centers of galaxies.

Easy experiment, hard analysis

Pulsar timing arrays monitor millisecond pulsars in the Milky Way. As a pulsar rotates, it emits radio waves that sweep by Earth with the period of rotation. The shorter the period, the more and sharper the incident radio-wave ticks, so the better the clock. A gravitational wave passing between a pulsar clock and an observer distorts spacetime and causes the signals to arrive either later or earlier. "That's the fingerprint," says EPTA member Alberto Sesana of the University of Birmingham.

The distortions due to gravitational waves are faint; if they arise from galaxy mergers, they occur with periods of decades. Data are collected for each pulsar for about a half hour every few weeks. The hundreds of thousands of pulses from a given observation are summed to extract the signal from the noise.

Extragalactic gravitational waves wash over all the pulsars in the Milky Way. Because the pulsars are independent, and each has its own timing and its own interstellar medium, the main giveaway for detecting a gravitational wave is a correlated signal between the pulsars. "Pulsars do their own thing, and it's hard to dig out a tiny signal from a large number of sources of noise," says Sesana. "The main way to overcome this is by timing an array of pulsars."

THE WESTERBORK SYNTHESIS RADIO TELESCOPE in the Netherlands brings 14 dishes, each 25 m in diameter, to the European Pulsar Timing Array project. Together, they have an effective diameter of 94 m.

Each experiment keeps tabs on around 50 pulsars. Often an individual astronomer is responsible for specific clocks. For example, NANOGrav member Maura McLaughlin of West Virginia University monitors five. Assessing the data “is not completely deterministic,” she says. “It’s a bit of an art form.”

A key complication is that gravitational waves are not the only source of correlations. “You need to mark the time of arrival exactly,” says Verbiest. To minimize timing errors, which can introduce false correlations, the observatories have their own hydrogen maser clocks. Correlations can also come from incorrectly estimating the positions of Earth and other planets in our solar system. “If the Earth is not where we expect it to be,” he says, “all pulsars are affected, but not all in the same way.” The location of the solar system’s center of mass varies by a couple hundred meters in different models, explains Siemens. At the sensitivity now reached by the timing arrays, those differences matter. “We have to account for the variations in our models,” he says.

Fortunately, because the correlations from timing errors, positioning errors, and gravitational waves take different forms—gravitational-wave correlations, which are quadrupolar, are the most complex—they can be disentangled.

Other complications arise from inhomogeneities in the interstellar medium, from pulsars slowing down as they lose energy, and from combining data from different observatories. The interstellar medium spreads out the signal in a wavelength-dependent way, such that high frequencies arrive before low frequencies, says McLaughlin. The interstellar medium introduces other noise too. “It’s important to remove those effects.”

Increasing sensitivity

The EPTA uses dishes at five observatories—the Radio Telescope Effelsberg in Germany, the Lovell Telescope in the UK, the Nançay Decimetric Radio Telescope in France, the Sardinia Radio Telescope in Italy, and the 14-dish Westerbork Synthesis Radio Telescope in the Netherlands. Combining raw data from the observatories before the analysis affords greater sensitivity than combining postanalysis. Although the dishes range from 64 m to 100 m in effective diameter, the combination “becomes comparable to Arecibo,” says Cees Bassa of the Netherlands Institute for Radio Astronomy. “It allows us to observe more and fainter pulsars.”

The 305-m-diameter Arecibo Observatory in Puerto Rico is the largest and therefore most sensitive of any radio telescope. (China’s Five-Hundred-Meter Aperture Spherical Radio Telescope, known as FAST, is not yet in full operation.) NANOGrav has won about 20% of observation time on it and 10% on the Robert C. Byrd Green Bank Telescope in West Virginia. The PPTA in Australia uses the venerable 64 m Parkes radio dish in New South Wales.

But Arecibo, Green Bank, and Parkes are on shaky financial ground and have been threatened with closure (see PHYSICS TODAY, February 2017, page 26). For now, Parkes has external funding. And the NANOGrav collaboration is looking for private donors. “We want to buy all the available time at Arecibo and Green Bank,” says Siemens. “That would be $12 million a year. It would save both telescopes and greatly increase our sensitivity.”

Looking ahead, pulsar timing array science has been identified as a key observational project for MeerKAT, which will see first light in South Africa next year. MeerKAT, FAST, and the Square Kilometre Array “will be a game changer in our business,” says Sesana. “We’ll be able to time pulsars better and to discover more pulsars. And with more pulsars, the sensitivity to [gravitational waves] gets better.”

In 2008 the three projects teamed up to form the IPTA. The first joint data set came out last year and contained a cumulative 559 years of data amassed from 49 pulsars. “It’s hard to quantify how much sensitivity you gain,” says Verbiest, who headed the effort. “A quick and dirty estimate says we gained by a factor of two.” The combining process was “a sobering experience,” he adds. “As much as we are a tight-knit community, each group still has its own methods and habits.”

Each team also wants to be able to make claims on its own, in part to increase its chances of getting funding. Sharing data is a sticking point, says Sesana. "When you share, you open the database of interesting physics to everybody. You might have students working on projects on specific pulsars, and you want to protect your students." Still, with scientists circulating among the projects and closer collaborations in the works for future facilities, says Verbiest, "the pulsar timing arrays are slowly growing to accept a more global approach."

Uncharted territory

The primary targets of the pulsar timing arrays are supermassive black hole binaries. "We think they form when galaxies merge," Sesana says. "We have circumstantial evidence and predictions, but the way to nail them and to understand them is by observing their gravitational waves." Scientists hope to learn how frequently they form, how many they are, how they evolve, and more.

The pulsar timing array experiments may also detect "exotic" sources such as erratic gravitational waves from when the universe was undergoing inflation and from superstrings—topological defects in the universe with huge concentrations of energy that could shake up spacetime with relativistic oscillations. In any case, Sesana says, "By looking at longer wavelengths than have previously been probed, we will unveil completely different phenomena."

Many in the community expect that the first detection won't be from a single event but from the stochastic background created by thousands of black hole mergers. "It would look like a long-period rumble," says McLaughlin.

The collaborations have not yet plucked a gravitational wave from their data. "But even nondetections place stringent constraints on the theory," says Verbiest. "The data have already been useful." There are a host of sources that might be detectable, he says. "The difficulty is that we don't know what to expect. We have to be careful as we get sensitive to things we are not aware of. We are probing uncharted territory."

Toni Feder

Cleanup of Cold War nuclear waste drags on

Despite billions of dollars spent preparing to treat and stabilize liquid radioactive wastes, cleaning out leaking tanks at the former nuclear production site in Hanford, Washington, will take decades more.

A quarter century after US nuclear weapons production ceased and cleanup of the sprawling Cold War weapons complex began, the US Department of Energy has achieved some notable successes. The Rocky Flats Plant in Colorado, for example, where plutonium pits for tens of thousands of nuclear weapons were manufactured, has been restored as a wildlife refuge. (See the article by David Clark, David Janecky, and Leonard Lane, PHYSICS TODAY, September 2006, page 34.) But stabilizing and safely disposing of the tens of millions of gallons of highly radioactive and toxic liquid wastes produced in making that plutonium remains a distant goal.

Although two DOE installations with liquid wastes have turned some of them into stable solids, the one with the largest inventory, the Hanford site in Washington State, has yet to treat any of the 212 million liters stored in its 177 underground tanks, some of which are about 16 kilometers from the Columbia River. Some of the waste dates to the 1950s, and leaks from at least 62 tanks have dumped an estimated 3.8 million liters into the soil.

Since the Hanford cleanup effort began in the early 1990s, DOE has spent $19 billion on several waste-treatment strategies, according to the Government Accountability Office (GAO). The target date for processing and disposing of all the waste by pouring it into canisters and converting it to a glassified state is now set for 2047, more than a century after the Fat Man bomb was dropped on Nagasaki. But even that timetable isn't definitive, since DOE has yet to specify how more than half the waste will be treated and disposed of.

Designed to turn the waste into glass logs—a process known as vitri fication (see the article by Ian Pegg, PHYSICS TODAY, February 2015, page 33)—Hanford's waste treatment plant (WTP) has been under construction since 2002. It was originally planned as a pilot plant to process 10% of the waste, and was estimated to cost $4.3 billion. DOE later decided to expand the project into a full-scale facility, and its cost to treat only a fraction of the waste has now ballooned to nearly $17 billion.

At the WTP, the tank waste is to be separated into a high-level waste (HLW) stream, roughly 10% of the tank volume containing more than 90% of the radioactivity; the remaining volume will be low-activity waste (LAW). A pretreatment plant will remove soluble fission products and return them to the tanks for treatment as HLW. Each waste stream will be mixed with borosilicate material to form a molten glass, which will be poured into stainless steel canisters. The HLW will be stored in an onsite vault for eventual disposal at a nuclear waste repository that doesn't yet exist. The LAW canisters will be permanently stored at Hanford.

Under a 2012 court order that was amended just last year, DOE must begin vitrifying LAW by the end of 2023. But the order doesn't require HLW glass-log production to begin until 2036.

Under pressure to meet the 2023 deadline, DOE elected in 2015 to build yet another WTP component that will bypass the main pretreatment facility for 76 million liters of LAW, filter out the solids and cesium-137, and feed the treated material to the vitrification plant. Completion is set for 2021. The current WTP cost estimate of $16.8 billion includes construction of the new facility and the LAW vitrification plant, plus the work that has been done so far on the pretreatment and HLW vitrification plants. It does not include completion of the pretreatment and HLW facilities, and