

Science powerhouses unite to help search for gravitational waves

5 December 2014, by Aaron Dubrow



Scientists from Caltech, Syracuse University and Abilene Christian University worked with computing experts at the Texas Advanced Computing Center to develop and test new methods for transmitting and analyzing data with an eye towards Advanced LIGO's future operations. Credit: LIGO

Scientists since Albert Einstein have believed that when major galactic events like supernova explosions or black hole mergers occur in the universe, they leave a trace. That trace, it is believed, takes the form of gravitational waves—ripples in the curvature of space-time that travel outward from the source.

For over a decade, scientists and engineers have engaged in one of the most ambitious research efforts ever undertaken: to design, build and operate the [Laser Interferometer Gravitational Observatory \(LIGO\)](#) in order to identify signs of gravitational waves.

From 2002 to 2010, hundreds of scientists worked together to bring the experiment to life. It required two gravitational wave observatories, located 1,865 miles apart and working in unison: the LIGO Livingston Observatory in Livingston, La., and the LIGO Hanford Observatory, located near Richland,

Wash.

LIGO was developed and is managed by the California Institute of Technology and MIT, and is funded by the National Science Foundation (NSF).

Since gravitational waves are expected to travel at the speed of light, the distance between the sites corresponds to a difference in arrival times of about ten milliseconds. By triangulating the signals, this small difference in detection can determine the source of the wave in the sky.

"Gravitational waves offer an unprecedented opportunity to see the universe from a new perspective, providing access to astrophysical insights that are available in no other way," said Stuart Anderson, a Caltech research manager for LIGO.

To date, no waves have been detected. Yet most astronomers believe they're still out there, waiting to be discovered with use of the right tools. For that reason, the National Science Foundation, the U.K. Science and Technology Facilities Council, the German Max Planck Society and the Australian Research Council have supported the creation of a follow-up to LIGO called Advanced LIGO.

Advanced LIGO replaces LIGO's original detectors with ones ten times more sensitive. It is expected to go online in 2015.


"Advanced LIGO's goal is to pioneer a new field of gravitational-wave astronomy, and its planned sensitivity will allow this very significant step beyond first detections," Anderson said. "The massive, relativistic constituents, causing violent perturbations of space-time itself, will be revealed through this unique probe."

Big Data and detection

In addition to all of the astronomical considerations

that must be attended to in a project like this—the placement of the observatories, the sensitivity of the detectors, the modeling of the expected signals—Advanced LIGO is also a "Big Data" problem.

The observatories take in huge volumes of data that then must be analyzed to determine their meaning. In the case of a positive identification of a gravitational wave candidate, the analysis must happen quickly, so the location of the waves can be identified and used for further study.



Advanced LIGO will take in huge volumes of data that must be analyzed to determine their meaning. Researchers estimated that ALIGO would generate more than 1 petabyte of data a year, the equivalent of 13.3 years of high-definition video. Credit: MIT/CalTech LIGO

Developing an offline system for binary neutron star and black hole detection was one of the many goals of the Advanced LIGO project, Anderson said. "We wanted to dig deeply into signals, since we have a good idea of what those signals are."

Since gravitational waves are still in their theoretical infancy, scientists require additional information to learn about their existence. Combining observations of electromagnetic radiation, neutrinos or [gravitational waves](#)—a process known as multi-messenger astronomy—provides complementary information, but it makes a messy data problem even messier.

When they were planning Advanced LIGO, the researchers realized they needed more computing and data processing power than they had used previously. But how much more did they need? And where would it come from?

By moving to more a sensitive detector, they had to increase their computing needs by more than an order of magnitude. They estimated that ALIGO would generate more than 1 petabyte of data a year, the equivalent of 13.3 years of high-definition video.

Discussing the matter with program officers at NSF, they decided to see if the [Extreme Science and Engineering Discovery Environment \(XSEDE\)](#) could help.

XSEDE, another NSF-supported project, oversees the nation's network of supercomputers and digital resources. These resources allow scientists to solve otherwise impossible computing problems and pursue really big science. Perhaps it could do the same for Advanced LIGO.

Scientists from Caltech, Syracuse University and Abilene Christian University worked with computing experts at the Texas Advanced Computing Center (TACC) to develop and test new methods for transmitting and analyzing data with an eye towards Advanced LIGO's future operations.

To achieve their goal, the team overlaid LIGO's computing environment on top of the [Stampede supercomputer](#) at TACC, one of the most powerful in the world. They then used a tool called Condor to break their large computing problems into smaller parts that could be distributed and solved on the supercomputer's individual nodes. (A node consists of a number of processors connected together and to the larger system.)

Miron Livney's High Throughput Condor group at the University of Wisconsin-Madison provided critical support (including software updates) to allow this computational challenge to work, and the members of the Advanced LIGO team were given priority access to the supercomputer so their jobs could run efficiently with few interruptions.

At first, they used a data analysis code from the initial LIGO for the experiment, since it was easy to validate. However over time, they began using the new Advanced LIGO version of the analysis code. Through detailed benchmarking, they were able to verify the reliability of the process and offer a preliminary sense of how much computing and data-crunching capability they needed to transmit, store and analyze data from the observatories for use by the astronomy community.



LIGO, the Laser Interferometer Gravitational-Wave Observatory, is a large-scale physics experiment that aims to directly detect gravitational waves. Here is the LIGO Livingston Observatory in Louisiana. Credit: MIT/CalTech LIGO

"The XSEDE and TACC teams were extremely helpful in collaborating on a successful port of a large scale LIGO data analysis pipeline to the Stampede system—all the way from network transfers of instrument data, through large scale number crunching, to generating and presenting web accessible scientific results," Anderson said.

"The solution we implemented allowed individual LIGO data analysts to use their existing computer identities and software tools to quickly and efficiently leverage the Stampede computing resources. I am very excited by the possibility of using XSEDE resources to maximize the scientific reach of the Advanced LIGO experiment."

Accelerating the Advanced LIGO code

Not only were the researchers able to prove they could compute on a big national supercomputer. In the process of collaborating with experts at TACC, they ended up making the Advanced LIGO software four times more efficient, ultimately saving money and time.

"We were able to reproduce the answers from previous experiments and also took advantage of our local experts to speed up the LIGO codes," said Yaakoub El Khamra, a research engineer at TACC who collaborated on the project. "The effort was an unmitigated rousing success."

According to Peter Couvares, a computer scientist from Syracuse University who worked on the project, the experience showed LIGO researchers that large high-performance computing systems like Stampede can work as well as the more familiar high-throughput computing systems that use a number of clusters distributed across LIGO Scientific Collaboration institutions.

"Given that our foremost computing challenge has always been to enable scientists to automatically and reliably execute large complex workflows consisting of millions of jobs and big data, NSF-supported tools like Condor and Pegasus have been invaluable," Couvares said. "We're increasingly realizing the value of HPC expertise at places like TACC to optimize our codes to make efficient use of shared XSEDE resources like Stampede."

Borne out in tests run over the last few months, the results proved that using the nation's supercomputers for gravitational wave analyses is feasible. In fact, in the end, the team was able to run more than 10,000 analysis jobs simultaneously on Stampede.

They presented their results at the HTCondor Week 2014 workshop in April.

"TACC staff not only got the LIGO team up and running on Stampede, they helped make the LIGO code substantially more efficient," said Dan Stanzione, TACC director. "Over the course of a year of computing with advanced LIGO, that could equate to millions of dollars in savings, or a huge

increase in the amount of processing we can do on the same budget. This shows the value of software engineering being interjected into the process."

Provided by National Science Foundation

APA citation: Science powerhouses unite to help search for gravitational waves (2014, December 5) retrieved 18 November 2019 from <https://phys.org/news/2014-12-science-powerhouses-gravitational.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.