tracking transients

how los alamos's thinking telescopes monitor and interpret millions of points of light

OUR UNDERSTANDING OF THE UNIVERSE and our place in it comes from observing the patterns and motion of all the lights that speckle the darkness. Even before the invention of the telescope, stargazers saw familiar shapes among the brightest points—a hunter, a scorpion, a ladle, a lion—and tracked the changing positions of the planets. But the first time a telescope was pointed into space, it heralded a new era of astronomy. With increased depth of vision, new objects were illuminated and new patterns revealed. Indeed, telescopes profoundly transformed our understanding of everything that fell under their gaze—at least, everything that sat still long enough to be seen.

A transient astronomical event is, unsurprisingly, one that lasts for a short time then is gone. Such transients can last anywhere from seconds to days and are caused by changes in the position or intensity of light. They occur when orbiting satellites sail by or distant cosmic explosions blaze into view, when aircraft navigation lights blink and meteors burn up. In a single night, millions of transients take place above our heads; the night sky is alive with momentary twinkles, flashes, and flares. Look up for a moment and you're bound to see one, and if you have a telescope, you'll see it even better.

Astronomy is observational by necessity. Until the not-too-distant past, it was a circadian science—go outside at night, watch the sky, come inside and record what you saw, repeat the process the next night, and see what has changed. But Los Alamos astrophysicist Tom Vestrand says the cadence of astronomy is changing. No longer constrained to a daily rhythm, developments in the night sky can now be observed minute-by-minute.

"We are entering an exciting new era of time-domain astronomy," he says, "where there will be an overwhelming number of transients found in real time." Time-domain astronomy means doing repeated scans of the sky then looking for changes from one scan to the next, and it's Vestrand's goal to get the time domain down to mere seconds. He is the lead on Los Alamos's RAPid Telescopes for Optical Response (RAPTOR), an array of "thinking" telescopes that are being trained to discriminate which transients to observe, independent from their human operators.

Thirty miles outside of Los Alamos, in the Jemez Mountains, lives the RAPTOR family of specialized telescopes. Vestrand likens the array to an ecosystem: multiple discrete elements serving different purposes and working together for an optimal outcome. They are the fastest-slewing telescopes in the world, able to swivel from point to point in under five seconds. There are wide-field lenses for persistent monitoring (like human peripheral vision) and narrow-field lenses that zoom in to take a better look at interesting events. The most recent additions are a 16-lens telescope that scans the entire sky every five minutes, and a telescope with four colored filters for enhanced observation of the optical emissions from celestial transients.

The RAPTOR robotic observatory is central to Los Alamos's Thinking Telescope Project, which is developing both the hardware necessary to track transients in real time and the software needed to manage the massive amount of data collected. The project already has robotic telescopes that can detect and identify transients automatically. Now Vestrand's team is training them in autonomy—teaching them to determine which transients are interesting and what type of observation to perform on a case-by-case, second-to-second basis.

Training telescopes

While the idea of thinking telescopes autonomously tracking millions of blinks and twinkles in the sky is thrilling, comparatively slow and clumsy humans can't be eliminated from the picture altogether. Central to robotic observatories like RAPTOR are high-powered computers that first need to be programmed to classify transients and initiate follow-up observation. Each time a transient is detected, it must be correctly classified according to its initially observed characteristics. What kind of event it is will inform what kind of follow-up observation is appropriate, and the follow-up data will feed back and help fine-tune the classification process.

gamma-ray bursts

A Los Alamos History of Discovery

In 1972 a Los Alamos scientist named Ian Strong was asked to dig through nearly 10 years of back-catalogued data collected by the Vela Satellite Program. A tiresome task to be sure, but it would lead to the discovery of one of the most curious modern astronomical phenomena and open the door to decades of fascinating and far-reaching research.

Vela's primary goal was to monitor the earth for illegal nuclear explosions following the Nuclear Test Ban Treaty of 1963. The Vela satellites were built at Los Alamos and configured to monitor x-ray, gamma-ray, and neutron emissions from the entire planet, 24 hours a day, 365 days a year. The gamma-ray detectors aboard the satellites were often triggered by events that were clearly unrelated to earthly nukes. Though interesting, these events were deemed peripheral to the program's main goal and were filed away for later study.

As he worked through the data collected from each trigger of Vela 5's gamma-ray detectors, Strong carefully calculated the source position in the sky for each one. There could be no doubt, he and his Los Alamos supervisor, Ray Klebesadel, finally reported, that the gamma-rays were of cosmic origin—meaning they weren't coming from the earth or its sun. This was the first evidence that major gamma-ray-emitting events were regularly occurring in deep space.

Prior to asking Strong to rummage through Vela's data, Klebesadel and another Los Alamos colleague, Roy Olsen, had noticed something odd. An event recorded by both Vela 3 and Vela 4 appeared to be a burst of cosmic gamma rays, but the instrumentation of the satellites was not yet sophisticated enough to pinpoint the source. After Strong's discovery in 1972, this earlier burst, which occurred on July 2, 1967, was retroactively designated the first-ever recorded gamma-ray burst (GRB).

But that was just the beginning. Many astronomers and astrophysicists weren't sold on the idea of the GRB's cosmic origin. They thought there were other plausible explanations. It wasn't until 30 years later, in 1997, that irrefutable proof in the form of direct observations of the optical afterglow of GRB 970228 finally snuffed the debate. The distances of subsequent GRBs placed them far outside our own galaxy, and the tremendous energy released suggested they came from colossal events, like when a star goes supernova or when two ultra-dense stars collide. Gradually, a clearer picture of what GRBs are, how they occur, and what they do was coming together—and it was breathtaking.

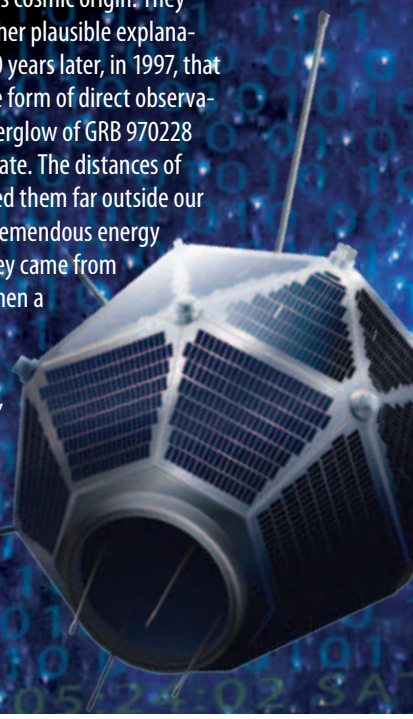

Przemek Wozniak, a scientist on Vestrand's team, is programming the computers to distinguish real transients from bogus ones. It's a balancing act between maximizing event detection and minimizing false positives. Transients are identified by the comparison of light position and intensity between adjacent time-domain scans of the same part of the night sky. The scans, taken roughly 5 minutes apart, are superimposed on one another and the first scan is digitally subtracted from the second. Any remaining light in the image represents a recent change in the position or intensity of a light source, and therefore represents a new potential transient.

A real transient is an astronomical event that occurs somewhere in space and is recorded. A bogus transient gets recorded when there is something that may look to the computer like an astronomical transient but isn't, like a stray cosmic-ray particle hitting the detector dead-on or an aircraft blinking as it passes through the telescope's field of view. Wozniak has created a dictionary of hundreds of real transients that the computer uses for comparison against candidate transients. By comparing his millions of candidate transients to the training dictionary, Wozniak identified a sweet spot between false positives and missed events and has been able to achieve a classification accuracy of 90 percent.

Once a transient event is determined to be real, follow-up observation begins, asking new questions and collecting new data. Much like an interview, the answer to each question informs the next question. For example, a car crash is an event. Was the driver speeding? No. Then was the road wet? Yes. Okay, so were the tires worn? No. Was the driver distracted? It is a real-time, iterative process designed to extract maximal, relevant data to fully describe a unique event. And since astronomers can't predict when an important event will occur, having the continuous monitoring system in place, with fast and accurate data acquisition capabilities, is key to capturing and understanding transient events in space.

The burst of the century

Gamma-ray bursts (GRBs) comprise one of the more dramatic cosmic events in RAPTOR's aim (See "Gamma-ray bursts: a Los Alamos History of Discovery" above). GRBs are extraordinary phenomena whose signals arrive on Earth

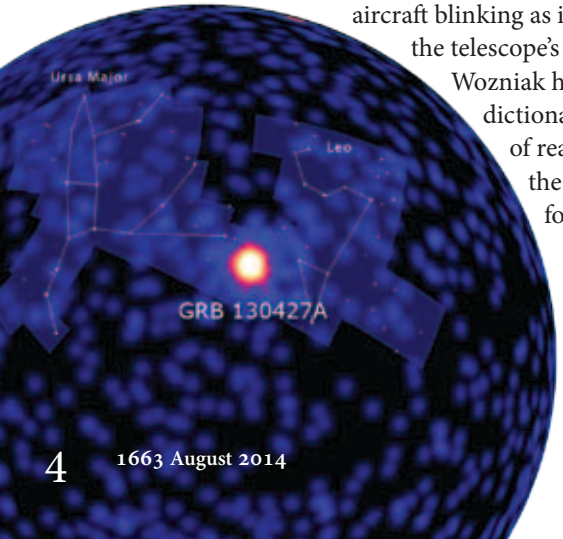

From in between the constellations Leo and Ursa Major (which includes the Big Dipper), GRB 130427A was one of the brightest, longest, and closest GRBs recorded by modern astronomical instruments, resulting in an unparalleled windfall of data.

Source: NASA

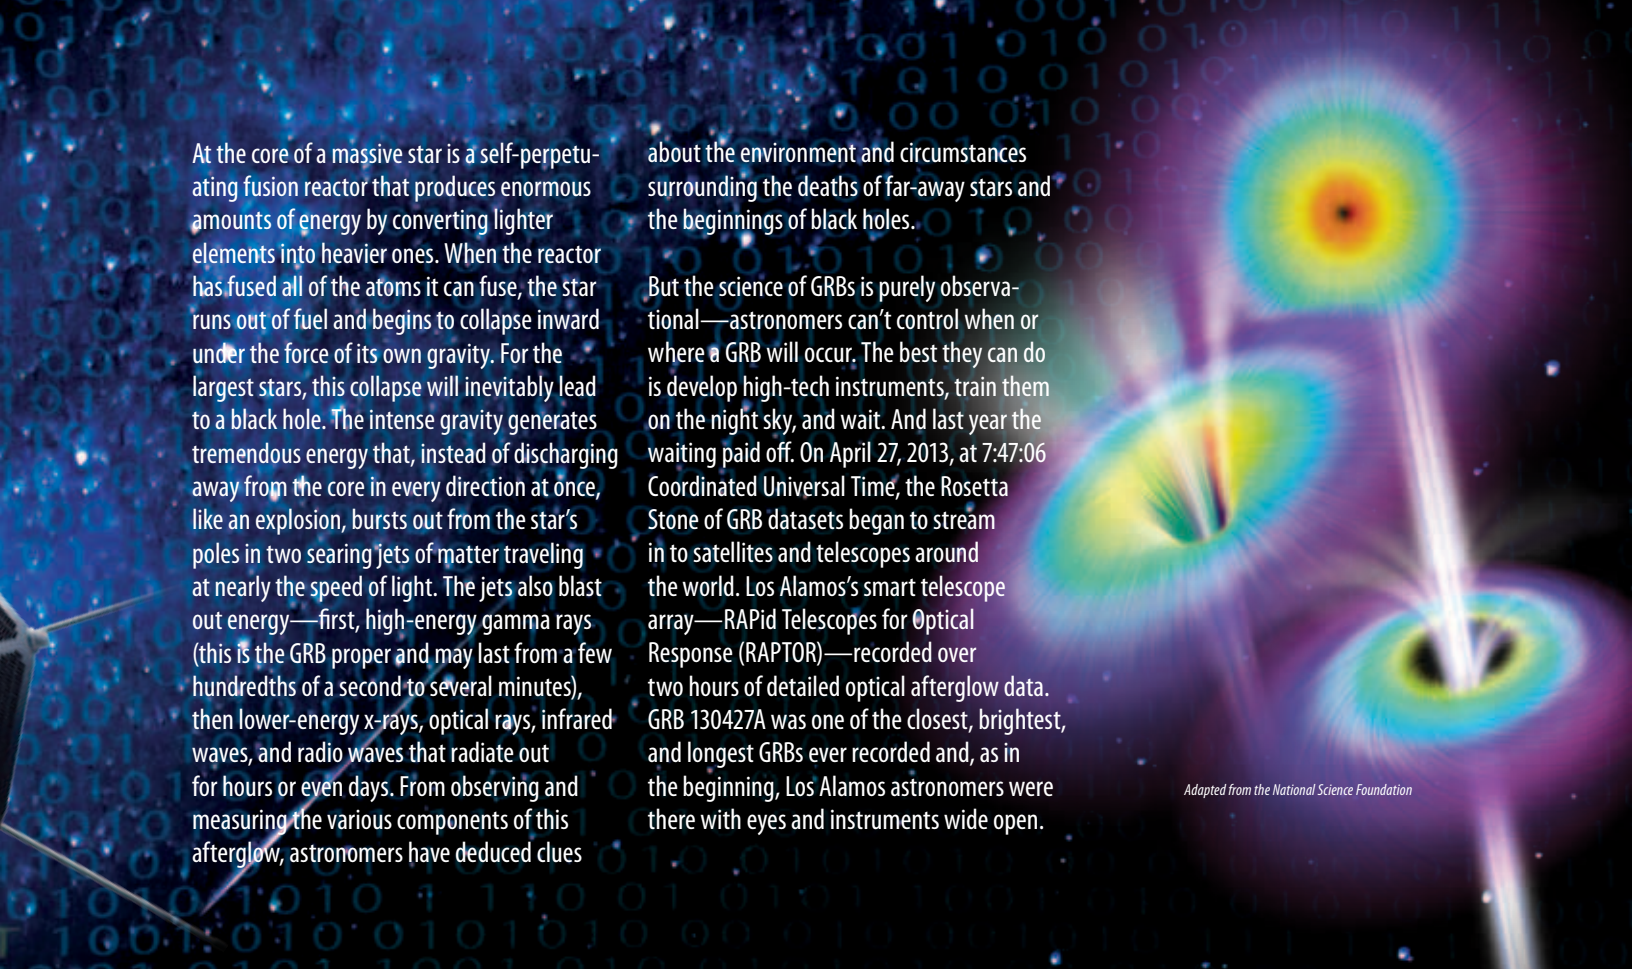

At the core of a massive star is a self-perpetuating fusion reactor that produces enormous amounts of energy by converting lighter elements into heavier ones. When the reactor has fused all of the atoms it can fuse, the star runs out of fuel and begins to collapse inward under the force of its own gravity. For the largest stars, this collapse will inevitably lead to a black hole. The intense gravity generates tremendous energy that, instead of discharging away from the core in every direction at once, like an explosion, bursts out from the star's poles in two searing jets of matter traveling at nearly the speed of light. The jets also blast out energy—first, high-energy gamma rays (this is the GRB proper and may last from a few hundredths of a second to several minutes), then lower-energy x-rays, optical rays, infrared waves, and radio waves that radiate out for hours or even days. From observing and measuring the various components of this afterglow, astronomers have deduced clues about the environment and circumstances surrounding the deaths of far-away stars and the beginnings of black holes.

But the science of GRBs is purely observational—astronomers can't control when or where a GRB will occur. The best they can do is develop high-tech instruments, train them on the night sky, and wait. And last year the waiting paid off. On April 27, 2013, at 7:47:06 Coordinated Universal Time, the Rosetta Stone of GRB datasets began to stream in to satellites and telescopes around the world. Los Alamos's smart telescope array—RAPid Telescopes for Optical Response (RAPTOR)—recorded over two hours of detailed optical afterglow data. GRB 130427A was one of the closest, brightest, and longest GRBs ever recorded and, as in the beginning, Los Alamos astronomers were there with eyes and instruments wide open.

Adapted from the National Science Foundation

daily but originate billions of light years away. To put that into perspective, consider this: our galaxy, the Milky Way, is approximately 100,000 light years across, so GRBs typically occur at a distance from Earth that would accommodate 10,000 Milky Ways strung end-to-end across the universe.

When a massive star in the distant universe has consumed all of the fuel available to it, it collapses in on itself. A black hole begins to form at the stellar core, and jets of matter burst from the stellar poles, producing high-energy gamma rays. If earthbound astronomers are fortunate enough to have one of the jets pointed toward Earth, they may detect it as a GRB. Then, following the gamma rays, which are light that is invisible to the eye, an afterglow of lower energy light, such as x-rays, optical (visible) light, infrared waves, and radio waves, continues to stream away from the moribund star. Most of what we know about GRBs comes from observation of the afterglow.

The GRB closest to Earth (3.8 billion light years) occurred in April of 2013 and had several superlative qualities beyond mere proximity. GRB 130427A, as it is known, was one of the longest ever recorded, with gamma emissions lasting over 20 hours (typical gamma emissions last only minutes or even seconds) and optical emissions lasting nearly two days. It was also one of the most energetic, which made it extremely bright. In fact, only one other GRB, in 2008, was brighter—so bright that its optical component could be seen from Earth without an instrument, earning it the nickname the "naked-eye burst." GRB 130427A was only slightly less bright than the naked-eye burst and was clearly visible through a standard pair of binoculars.

What made GRB 130427A a keystone event was not only that it was the closest, longest, and second-brightest GRB ever recorded, but that it was recorded in its entirety by 58 telescopes across four international collaborations. As Wozniak put it, "You need a very well-observed event with data from a wide variety of instruments, as much data as possible, to build a good GRB model. The April GRB was just such an event."

Another remarkable thing about the data collected by RAPTOR that night is that it begins slightly *before* the GRB itself. Because the system is set up for persistent monitoring, looking everywhere all the time, it doesn't need to be told from an outside trigger that something interesting is happening somewhere. It's like how a witness to the car crash might recall the events leading up it—she didn't know she was seeing pre-crash activity until the crash began. RAPTOR's wide-field scopes were watching, and when the computer noticed an anomaly (which turned out to be the

early optical emissions), it automatically slewed its deep-field scope to a point in between Leo and the Big Dipper to get a better look—all without waking Vestrand, Wozniak, or any other human for guidance.

With a landmark dataset from an unprecedented GRB event, Vestrand and collaborators now challenge current theoretical understanding of how GRBs work. They found gamma-ray photons (particles of light) with calculated energies at impossible levels, up to 95 billion electron-volts (95 GeV). According to current GRB models, there is no way to produce photons with that much energy. And yet there they were. The data and interpretation were corroborated by independent analyses by NASA's Fermi Gamma-ray Space Telescope, Swift Satellite, and NuSTAR telescope. So, if the models can't account for these photons, and the photons are undeniable, then the models, it would seem, have to change.

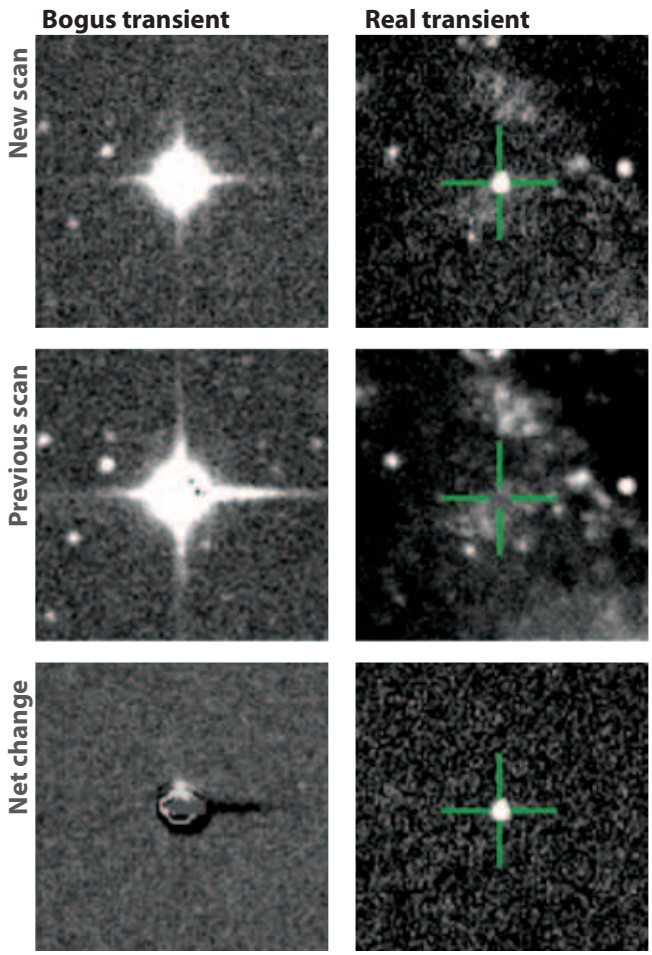

The computer vision of RAPTOR can automatically distinguish between a real transient and a bogus one with 90 percent accuracy. Shown here are a real transient (right column of images) and a bogus one (left column); in each, the top and center frames were taken at different times and are subtracted to show what has changed (bottom frame). By matching observations of the net change to a database of known real and bogus transients, a probabilistic score is calculated to indicate the likelihood of an event being real or bogus. The real events are then candidates for extended follow-up observation.

There exists a suite of standard proposed models for how GRBs operate, and a lot of data is required to fit any one model. The previously accepted model basically states that a generic GRB involves two distinct shock waves: a forward internal shock followed by a reverse external shock. It was generally thought that the gamma-ray emissions came exclusively from the internal shock and carried million-electron-volt (MeV) energies. However there was a smattering of evidence that a few gamma-ray photons, the highest energy ones, seemed to be delayed relative to the majority. This was a bit of a puzzle. If all of the gamma-ray emissions come from the same source, the forward internal shock, then they ought to have similar timing and energy. GRB 130427A cracked the puzzle apart. The few GeV gamma-ray photons were observed to correspond in time with the peak of optical emissions, which came from the reverse external shock. Therefore, the GeV gamma-ray photons came from the reverse external shock, not the forward internal shock, as stated by the previously favored model.

Such strong validation of a model for events that occur in the far universe is exceedingly rare, so GRB130427A is now a superstar among GRBs. But this was not the first chance for RAPTOR to prove its mettle. Since its inception in 2001, RAPTOR has observed countless GRBs and has been a key player in numerous international robotic astronomy collaborations. RAPTOR's observations of the optical counterparts of cosmic events have helped define and test theoretical models of GRBs, supernovae, and other cosmic transients for the past 10 years.

From generalists to specialists…

The ecosystem analogy that Vestrand uses to describe RAPTOR is apt: heterogeneity in instrument capability and autonomous interaction between the various elements are essential. Now that the instruments are in place, his efforts are focused on the intercommunication aspect. He and Wozniak are working on what they refer to as a dynamic coalition architecture that will reduce redundancy of observation and improve efficiency of data management.

"Right now, transient follow-up takes a second-grade soccer approach," says Vestrand. "All the players cluster around the ball and kick at it. The trigger comes in and all the telescopes of the world slew over and observe in some way." The dynamic coalition architecture, in contrast, will allow cooperation between RAPTOR and other robotic observatories and will provide customized task delegation for each transient that comes along.

In addition to training the RAPTOR telescopes, Wozniak is the primary investigator on a Los Alamos multi-institution collaboration called the intermediate Palomar

The RAPTOR robotic observatory is an ecosystem of telescopes. (Left) RAPTOR-Q (Queue), affectionately known as RQD2, is the wide-field element, providing the peripheral vision for the system with four bottle-cap sized lenses that together can see the entire night sky. It is pictured, left to right, with James Wren and Przemek Wozniak. (Center) RAPTOR-T (Technicolor) is one of the newest additions and consists of four identical lenses, each with its own color filter specially designed for observing the optical counterparts of celestial transients. Left to right are Cade Hermeling, Alin Panaitescu, James Wren, Tom Vestrand, Przemek Wozniak, and Heath Davis. (Right) RAPTOR-Z (Zippy) is the world's fastest telescope and provides the deep-field, up-close vision of the array. Next to RAPTOR-Z is team leader Tom Vestrand.

Transient Factory (iPTF). The iPTF uses a camera to systematically explore the transient sky and has emerged as the leading proof-of-principle experiment to an enormous and much-heralded new undertaking called the Large Synoptic Survey Telescope (LSST). The LSST will provide an exceptionally wide-field survey of the entire sky every few nights and will be the widest, fastest, deepest eye of the new digital age. Among other impressive functions, it will collect data on over two million transients per night. That's about 55 events per *second*. This is where the dynamic coalition architecture comes in. It will have an event broker function, so each time a transient is detected, the broker will consult every instrument in the network, essentially saying, "I've got an event here, and based on what you said you're looking for, I recommend that you take a look at it." In this way, each observatory can specialize in its follow-up and the collective data will be complementary and complete.

But there are some old-fashioned challenges to the new-fangled future of sky watching—international collaboration, for instance. The more observatories that enroll in the coalition, or subscribe to the event broker, the more efficient and useful it will be. Yet astronomers tend to be very protective of their machines, and understandably so. "Most people don't want to give me the rights to take over their telescopes," Vestrand points out. Each institution and researcher has time, resources, reputations, and ambitions invested in their programs. How do they weigh the greater good when their careers are on the line? So the coalition manager needs to know and cater to the priorities of each telescope, which comes from their human owners. It's an iterative process though, with each round of refinement needing extensive field-testing, which takes time and presents another challenge to the humans involved—patience.

…and back to generalists

The persistent-monitoring and data-interview model also has applications to fields other than astronomy. "Once you solve it at an abstract level you can apply it to any set of targets and assets you choose," Wozniak says. The technology is transferable—in any instance where it's important to maintain automated awareness of potentially dangerous patterns or anomalies buried in large amounts of noise, a dynamic coalition architecture can be applied. Biosecurity is an immediate example that Vestrand and Wozniak are not working on but could eventually benefit from the algorithms they are developing. Identifying how an infectious disease outbreak began—whether natural or nefarious—and predicting and mitigating the risks associated with it requires first sifting through incidental chatter to find the pattern and then projecting the pattern into the future. By organizing and interpreting that chatter in a global dynamic coalition architecture, perilous trends might be spotted sooner and dire outcomes avoided.

But the problem is mostly front-end loaded. Once the heavy lifting is done—the programing, training, validating, networking, and streamlining—then the sky is the limit for what sorts of questions might be answered and what sorts of discoveries might be made. Astronomy, it seems, has burst open like a supernova and will never again be the slow and lonely endeavor it once was. LDRD

—Eleanor Hutterer