JAMES L. OLDS

It's a Small World (Network) After All

The success of the collaborative effort to detect gravitational waves holds lessons for how to strengthen the resilience of the US science ecosystem.

One point three billion years ago, two stellar-mass black holes collided with one another with cataclysmic results. The collision produced a merged celestial body and, as a result, warped the universe's fabric of space-time in much the same manner as a pebble thrown into a pond produces ripples. The ripples from that merger traveled across space to Earth and, on September 14, 2015, were detected by warehouse-sized measurement machines built by humans here in the United States. To be believable to researchers, there needed to be more than one machine; in fact, two were built, and each had to "hear" an identical ripple in space-time.

These machines were so precise that they could reliably detect a movement in space of less than 1/10,000th the diameter of a proton. When the black holes merged, their giant masses resulted in a tiny physical movement in the detector (something like a very precise measuring stick) here on Earth. The machines' laser sensors responded to the movement, and high-performance compute flagged the movement as the marker of the collision so far away and so long ago. Not only did the two machines verify the other's result, but the seven millisecond delay between the two "chirps" also localized the collision to a particular quadrant of the sky. That discovery ultimately resulted in a Nobel Prize in Physics in 2017.

But this is not a science story—it's a social network story. One of the detection machines was in Hanford, Washington. The other was in Livingston, Louisiana. The two widely separated detector sites represented decades of distributed collaboration between multiple institutions, all funded on a billion-dollar long-shot bet by the US National Science Foundation. The September 2015 success was "imagined" 30 years earlier by people willing to take big risks—not only with their careers, but also with taxpayer dollars.

The success of this project, the Laser Interferometer Gravitational-Wave Observatory (LIGO), depended on a specific network topology, not only of the two remarkable devices, but of scientists, institutions, program officers, and the US government (including both the executive branch and Congress). To use the technical language of network scientists, they were linked by social, professional, institutional, and financial "edges" and represented key "nodes." As with the two machines, one research lab alone was not sufficient for the job at hand.

The discovery of gravitational waves, predicted by Einstein but not "heard" until that September 2015 moment, was an emergent property of the US science ecosystem, the same one that science administrator Vannevar Bush described in *Science, the Endless Frontier* in 1945. In the years since, that enterprise, which has been called a "miracle machine," has produced such wonders as the transistor, GPS, the internet, and eventually the transformer models that power artificial intelligence.

In his report to the president, Bush proposed a compact: The federal government would fund the basic research that the private sector would not fund because of profitability concerns; such research would benefit the nation's economy and overall health. The example of early-twentieth-century discoveries in physics, such as quantum mechanics leading to the atomic bomb, was very much on the author's mind.

To my mind, LIGO is the clearest possible test case for the Bush compact of federal patronage of basic research with no

requirement to justify itself in market terms. LIGO was neither applied nor practical science. There was no product, no patent, and no near-term economic case.

How LIGO stayed alive

Even with the compact in place, LIGO survived many brushes with death. In 1991, LIGO received its first-year funding of $23 million from the Mathematical and Physical Sciences Directorate of NSF; however, a significant contract was temporarily frozen in 1993 over management concerns. In 1994, those concerns were addressed through a restructuring under a new principal investigator, California Institute of Technology's Barry Barish. The decisive move occurred shortly thereafter when NSF secured buy-in from legislators representing the host state of Louisiana and from Massachusetts (the Massachusetts Institute of Technology partnered with Caltech on LIGO). The project had become a geographically distributed coalition. Some observers could call that pork barrel politics, but such politics are also governed by the scientific rules of networks. With nearly $400 million in additional funding approved in 1994, LIGO became NSF's largest project to that point.

When the initial experimental setup proved unsuccessful in observing gravitational waves, NSF went double-or-nothing with Advanced LIGO. The cumulative cost was now more than $1 billion. But when those machines were turned on in 2015, the black hole merger was observed immediately after. The device opened the window to a new kind of "multi-messenger astronomy" combining LIGO with Earth-based optical and radio telescopes along with space-based platforms like the James Webb Telescope. Most recently, during the 2025–2026 congressional appropriations fight, LIGO funding was defended against proposed executive branch cuts. The pushback was bipartisan.

So why did a project with zero practical payoff survive nearly three decades, multiple administrations, and a near-death experience within NSF? I've had a personal role in a similar near-crash of an NSF megaproject, the National Ecological Observatory Network, which got into trouble while I was heading up NSF's Biological Sciences Directorate. So I speak with some level of understanding.

First, of course, the science underlying LIGO was excellent. Gravitational waves lie at the very basis of Einstein's theory of general relativity. Observing them fundamentally deepens scientists' knowledge of the universe. But the deeper political answer is that LIGO was embedded: It was connected to multiple institutions, states, congressional delegations, and scientific disciplines (physics, engineering, computer science, astronomy). What if that embeddedness isn't incidental, but the actual load-bearing structure of the whole American science ecosystem?

In other words, the architecture Bush proposed is not what made American science so successful. Although that framework was largely instantiated in 1950 when Congress passed the law creating NSF, it was the emergent network that followed on as a secondary effect that serendipitously resulted in the system's fantastic productivity and agility. Bush's description of how science and engineering would deliver public benefits to the US public was largely linear, but it was the ecosystem's nonlinear network that constituted the true miracle machine. This connected network provides explanatory power for major successes, including LIGO. And this is the fantastically productive entity that China has struggled to copy.

The power of the network extends beyond the LIGO example. Consider the discovery and development of CRISPR, the gene-editing technology that revolutionized molecular biomedical research. While LIGO's goal was to test Einstein's theory of general relativity, CRISPR's success was far more serendipitous. It derived from the study of small palindromic repeats in bacterial genomes that were the signature of a primitive immune system. It was the network of other researchers who recognized how that system could be engineered into a practical tool that was momentous. In LIGO, the system worked out the engineering to test a theory. In CRISPR research, the network took an unexpected experimental result to engineer a novel tool. Both successes resulted from the emergent qualities of embedded networks.

A very big small-world network

We all intuitively understand small-world networks from the idea of "six degrees of separation," which places each of us only six social connections away from any other person (like Kevin Bacon, for instance). The small-world concept was developed by the social psychologists Jeffrey Travers and Stanley Milgram in 1967, building on an idea from a 1929 short story by Frigyes Karinthy. Computational social scientist Duncan Watts and mathematician Steven Strogatz formalized the idea as a network property in 1998. In small-world networks, most nodes (scientists, funders, champions) are not directly connected to each other, but any node can reach any other in a small number of steps. The key mechanism for small-world networks is a mix of tight local clustering (dense subcommunities) plus a small number of long-range "shortcut" links that bind the subcommunities together. Beyond the canonical Kevin Bacon example, one can see these networks in airline hub-and-spoke systems, the power grid, and the internet's backbone.

Now, consider the topology of the US science ecosystem. The governmental funding hubs are self-evident: NSF, the National Institutes of Health, the Departments of Defense and Energy, NASA, and the US Department of Agriculture are the main ones. There are 187 research universities classified as R1 institutions (very high research spending and doctorate production) and nearly as many academic medical centers, as well as 17 national labs and various other types of entities that serve as the primary loci for individual principal investigators. Peer review study sections and panels, professional societies, journals, and preprint servers are the small-world shortcuts that

let a discovery in one cluster reach and influence a distant one quickly. Multiple levels of governance and flows of money and people further help the network to cohere.

The US science ecosystem has many of the characteristics of a small-world network. First, there is high local redundancy: If one connection fails, neighbors have alternate paths. The technical term for this is a high clustering coefficient. Second, the average path length in a small-world network is short. This means that information and resources move very quickly across the whole system. Great examples of these shortcuts are the rotator executives (called IPAs after the Intergovernmental Personnel Act that covers them) who serve as program directors and, in the past, also headed directorates at NSF (directorates are now led by senior government executives). In my role heading the Biological Sciences Directorate, my connections to the community through my continued embedding as a professor of neuroscience at George Mason University added crucial cross-pollination to life sciences investments during that period. The current administration ended that particular shortcut.

The redundancy of this ecosystem is also pretty clear. Multiple agencies can fund adjacent work (e.g., a biological question might be fundable at NSF, NIH, or DOE); multi-institution collaborations like LIGO spread risk geographically and politically. Congressional delegations also work as an additional, often-overlooked layer of the topology: Grants create political shortcut ties between institutions and members of Congress across party and state lines.

But the key property of networks like the US science ecosystem is one that canonical small-world networks do not formally require: degree-heterogeneity. Its links are distributed with extreme unevenness. This is the property that matters the most. Only a half dozen main funding hubs serve thousands of institutions. NSF alone maintains active funding relationships with over 1,500 institutions; a typical principal investigator maintains a handful. That's a difference of 2–3 orders of magnitude.

Networks with these three properties can be described as "robust yet fragile." They tolerate random node failure remarkably well; lose an average node and the network barely notices. The system is optimized against random shocks (e.g., single principal investigator misconduct, university scandal, a year of small funding cuts, the petering out of one kind of research). But when a highly connected node is lost, the whole system becomes acutely vulnerable, as has become obvious over the last two years.

These characteristics help explain why the US science ecosystem survived 80 years of administrations, recessions, Sputnik-style panics, and the end of Cold War funding. The network absorbed individual shocks locally without cascading. This resilience helped LIGO survive its near-death experience in 1993. But for networks with this specific set of characteristics, the same topology of nodes and connections that provides resilience to local shocks also creates big vulnerabilities.

An ungraceful fragmentation

Two documents stand out in articulating the current administration's science policy: Project 2025, the Heritage Foundation's federal policy agenda, and the recent *Science: A New Golden Age* report from the Office of Science and Technology Policy. Although neither document explicitly uses the language of network science, their recommendations for change target some of the most vulnerable hubs and connecting links of this network. This includes the governance of agencies like NSF and NIH, agency-level spending authority, indirect cost support, and the special historical role of US research universities.

These documents and other actions by the current administration are often framed in terms of making the country's science system more productive. For instance, the *Golden Age* report notes that "productivity has slowed," then explicitly recommends emphasizing the funding of individual scientists over "legacy institutions." As another example, a recent memo from acting NSF director Brian Stone describes a "concrete ambition" of doubling the productivity and impact of federal research dollars in the next decade before describing a suite of changes, including "investigator-driven funding." The idea of funding people rather than institutions is not new. It has been a hallmark of organizations such as the Howard Hughes Medical Institute, among others. Still, policy changes at the scale suggested by the current administration risk atomizing successful principal investigators into disconnected—but funded—bubbles while fragmenting the network.

Already, significant emergent effects are occurring across the existing network. For example, capping indirect cost reimbursements will harm the institutional hub layer (the R1 universities and academic medical centers that perform most federally funded basic research) by removing the mechanism that cross-subsidizes shared scientific infrastructure. Dismantling the National Science Board as an NSF governance and oversight structure, along with the many science advisory committees across the system, weakens the linkages that help set long-range direction and insulate agency decisions from short-term political cycles. Grant terminations targeting specific fields and institutions weaken local clusters directly. And proposals to change peer review affect the shortcut mechanism that lets expertise travel fast across the network.

Network science predicts that systems like the United States' science ecosystem don't fail gracefully. They fragment into isolated clusters, path lengths spike up, and previously close nodes (e.g., institutions and principal investigators) lose fast access to one another. The majority of the system may be untouched, and line items in budgets may survive, but the productivity of the whole falls as it becomes more fragile.

Further emergent symptoms of this fragmentation are becoming visible, including a brain drain of American scientists abroad, slower cross-institutional collaboration, and the loss of the political shortcut ties that historically built bipartisan

coalitions to support science projects. Such fragmentation changes the shape of what science gets done. Small, locally findable, low-collaboration work may survive, while LIGO-scale projects will fall by the wayside. Without the miracle network, the everyday innovations that constitute science's social contract and make Americans' lives better may also be at risk. Without the network, would iPhones have been built from thousands of academic research projects? Would there be millions of people who have lived longer after diagnosis with once-deadly cancers because they took checkpoint inhibitors? These innovations were not single shots from single scientists, but the product of transnational networks where, for example, the discovery of oncogenes could spur an understanding of kinase-cascades that was engineered into a class of effective drugs over decades.

So much has been dismantled in the former system that returning to the status quo ante is unrealistic. And despite its many strengths, the US science ecosystem had real problems, as I've written about elsewhere. Looking to the future, we in the US science enterprise must take the lessons of the network and its current fragmentation into account to build something better and more resilient.

Rebuilding science's network architecture

If the miracle in the machine is the network, then we need a completely new roadmap to protect it. Rather than understanding it through linear benchmarks such as budget line items, star scientists, or individual institutions, we must prioritize the health of the connected web of hubs and links.

To my mind, the key to building this roadmap is to view the science ecosystem like a master gardener, much as landscape architect Frederick Law Olmsted would have regarded the space now occupied by Central Park in New York before it was a park: He saw the constraints of geology and soil and then, rather than dictating that an oak tree or daffodil plot be in a particular location, he imagined the entire future park as a whole. How can we plan over the long haul, as Olmsted imagined a city park that would thrive for a century or more in the future? Here, network science can guide us.

Some lessons are prescriptive: A future social network architecture for American science should include explicit hub redundancy so that no single mechanism (one board, one agency) is a single point of failure. The load-bearing structure for the ecosystem should be widely distributed, as LIGO distributed its own political base across states and disciplines.

Next, we must work to protect and develop the shortcuts, not just the nodes. These include preprint servers, peer review, cross-agency cofunding, shared resources, multistate collaboration, and IPA rotators. These edges are as important as the nodes themselves. To be useful, and to achieve those sought-after impacts, scientific knowledge must flow as freely as possible, both inside and outside the country.

At the same time, the people and institutions who make up the network must learn to think of themselves as a whole, rather than as atomized competitors. As the network has been challenged politically and financially, the science community has largely decided to "duck and cover," mimicking the nonsensical advice given in classroom films in the 1960s about how to survive nuclear war. This is not an effective strategy for individuals, and it is already enabling fragmentation of the whole.

Instead, the community should learn from past efforts at organizing itself. For example, in 1975, when the potential hazards of recombinant DNA research started to become clear, researchers and others organized the Asilomar Conference to create voluntary standards. Much more recently, when it became clear that this administration and Congress were likely to push for changes to indirect cost reimbursement, institutions, associations, and individuals came together to create the FAIR model. These are examples of bottom-up mechanisms in which the community used the network itself to alter its emergent trajectory for the better.

And this is not merely about network as science. What Asilomar and FAIR also represent are politics. Science, ultimately, is going to have to learn to do politics in a new way. Going beyond the usual consensus budget line-item "ask," we scientists need to do the internal politics to get our own house in order, along with the external work of building bipartisan efforts for funding. This will involve reimagining the network's topology and then making change at every level—from individual labs to graduate students to scientific societies. Taking ownership and responsibility for the network is very different from the community's business-as-usual approach.

Members of this network also need to learn from previous systemic problems, such as the practice of gaming the system for short-term gain. For example, tuition has become increasingly important for research institutions' business models. As tuition goes up for doctoral students in science, technology, engineering, and math fields, so does the temptation to subsidize those costs through grant budgets, which quickly becomes unsustainable. This harms PhD students, as well as departments and institutions, and dilutes the impact of research funding. Gaming is just one of many dysfunctional practices that the system has permitted, sometimes over decades. The network's nodes and hubs must develop mechanisms to support the functioning of the whole.

This is what I mean when I say that those of us who are components of this network must stop seeing ourselves as independent agents and instead take a more global "gardener's view" of what we have all built together. Doing that may lead to a new compact with the American polity—one that prioritizes the need to cultivate (and curate) the network.

James L. Olds *served as assistant director of the National Science Foundation Directorate for Biological Sciences from 2014–2018 and is university professor of neuroscience and public policy at the Schar School of Policy and Government at George Mason University.*