INFRASTRUCTURE SIMULATION EFFORT HAS HIGH HOPES, FACES HIGH HURDLES

By Greg Goth

The continuing availability of vital resources is paramount in ensuring the US national economy's well being. Yet until recently, no attempt was made to form an overarching computational architecture capable of helping policymakers identify and protect those resources on a national level.

The National Infrastructure Simulation and Analysis Center, operated by a core partnership at Sandia and Los Alamos National Laboratories, is being envisioned as the central component of this protective architecture. Although the concept of NISAC originated prior to the terrorist attacks of 11 September, it was only in the attacks' aftermath that government officials and simulation experts recognized the center as a potential keystone in national defense.

NISAC has received $20 million in the 2002 fiscal year and should be fully functional by the 2005 fiscal year, with advanced modeling and simulation capabilities providing policy analysis, mitigation planning, education and training support, and real-time crisis assistance for a wide variety of users. It also received mention in post-attack legislation—the *USA Patriot Act*—"to serve as a source of national competence to address critical infrastructure protection and continuity through support for activities related to counterterrorism, threat assessment, and risk mitigation."

"The concept originated two years ago," says Lillian Snyder, NISAC technical lead at Sandia. "The people at Los Alamos and ourselves tried very hard to kick this program off and received small initial funding to flesh out the concept. Last year, we called together the leaders of government agencies and asked them, 'Does a capability like NISAC currently exist, and if it doesn't, is it needed?' They all said 'It doesn't, and it is needed.'"

The laboratories plan to leverage their modeling and simulation experience in areas as varied as nuclear weapons and traffic patterns, as well as a US$160 million infrastructure modeling and simulation development investment, to establish NISAC.

"We have the requirement to be able to assess the national enterprise," says Chris Barrett, the Los Alamos NISAC technical lead and leader of basic and applied simulation science at the lab. "We have to be able to get our arms around this interdependent sequence of things trickling through infrastructure to infrastructure, sector to sector, and that's the goal of the analytical capability NISAC will eventually provide."

Starting from not quite scratch

Getting its arms around the national enterprise will tax the most advanced technology the laboratories possess, but before it can even begin modeling scenarios, political and philosophical differences among many constituent groups will have to be negotiated. The definition of what critical infrastructure is will need to be focused, as will who is allowed to make those definitions and use that information.

The relative unavailability of terrorism response scenarios also factors into the early obstacles facing NISAC.

"I haven't found any past resources or reference material, as far as public literature from university professors and defense researchers," says Roger Smith, vice president and group chief technology officer for Titan Systems and area editor for distributed simulation for *ACM Transactions on Modeling and Computer Simulation*. "Since September 11, they're just beginning to put stuff out."

Smith has written a paper on the topic since the attacks, "Counter Terrorism Simulation: A New Breed of Federa-

1521-9615/02/$17.00 © 2002 IEEE

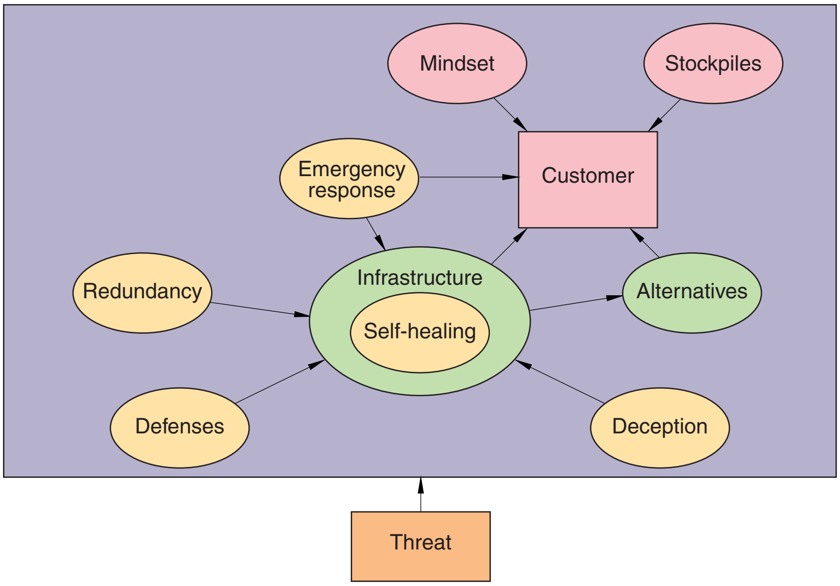

Figure 1. Infrastructure model that uses algorithms as nodes.

tion," and is completing a draft of another, "Complexities of Simulating Domestic Infrastructure Protection," which will be delivered in September.

"As soon as you stick one toe in this area, you realize how huge the term *infrastructure* is," he says. "You just can't grasp it. So, when you start subdividing it into power, water, telephone networks, and the other categories, you start getting something you can manage.

"If you just start building something, you're going to end up with the Great Wall of China. It'll be huge, unwieldy, and once you get it to work, you can't toy with it and change it because you'll break it."

One of Smith's potential infrastructure models involves nodes represented as algorithms that respond to stimuli from various sources. Each node would represent the infrastructure and its native ability for self-healing.

Los Alamos's Barrett says the nuts and bolts of NISAC will depend on how the Office of Homeland Security is organized and how it develops a clear mission statement. Much of the work this past year went into defining problems and possible solutions in workshops across multiple government agencies.

"If we're going to build decision-making tools to inform anyone, they have to be responsive to whatever this organization is," Barrett says. "Since it doesn't exist exactly, we're part of its creation."

Barrett acknowledges the political difficulties in convincing people to share information, but also contends that the problem is not new with NISAC.

"The ISACs [Information Sharing and Analysis Centers] culled out of *Presidential Decision Directive 63* [the 1998 executive directive on establishing critical infrastructure protection mechanisms] were a good example of just how hard it is to work through some of these problems. There are lots of people who feel we should be talking to them—hundreds of thousands of people. There's going to be a tension, especially now when things are not particularly well formed."

Example architectures

Barrett says the inability to access certain data might dictate the simulation modeling philosophy at NISAC.

"We have tried to use the most advanced techniques to get a handle on how to do unprecedented infrastructure simulations and cover questions that either haven't been answered before or that require data we can't get access to," he says. "The simulation techniques are heavily bottom-up. If you know how piece parts work, and you have some knowledge of process, you can substitute a lot of data and put these things together."

Perhaps the best known of the existing bottom-up simulations is TRANSIMS (the Transportation Analysis and Simulation System), a multimodule system that breaks metropolitan traffic patterns down to individuals within households and vehicles, and physical features such as parking lots, roads, and bridges.

TRANSIMS's modules (Selector/Iteration Database, Population Synthesizer, Activity Generator, Route Planner, and Traffic Microsimulator) reside in a flexible framework, which also includes data input protocols and defines intermodule dataflow.

The "glue" of TRANSIMS is the Selector/Iteration Database, which controls the flow of information among modules in methods as simple as feed-forward through the primary modules, or as complex as feedback loops that take data and feed it back to the Activity Generator and Route Planner modules for iterative computation.

The hardware for such modeling, Barrett says, evolved from CM-2 and CM-5 architectures to ASCI-style architectures, and into multiprocessor Linux clusters.

"We've produced an architecture and software implementation that rides on top of advances in SMP [symmetric multiprocessing] and MPP [massively parallel processing] architectures," he says.

In TRANSIMS studies in Portland, the architecture used

Computing in SCIENCE & ENGINEERING

consisted of Linux clusters with 64 dual-processor nodes with Intel Pentium III 500-MHz processors, 1 Gbyte of memory per node, a 100-mbit Ethernet connection between the nodes, a RedHat Linux 6.2 operating system (Linux 2.2.16 kernel), and 100-Gbyte disk accessible from all nodes via NFS. Another architecture used in the Portland study was a Sun Microsystems Enterprise 4000 with 14 Ultra Sparc II processors, 248 MHz, 6 Gbytes of memory, and 300 Gbytes of disk on a Solaris 7 operating system.

These architectures were designed to use the power of parallel computation to decrease the module execution time and leverage modern computing systems. Barrett says the decisions about architecture also take into account the tools the research is expected to produce for wider use.

"We don't feel constrained to only use what's available," Barrett says. "In fact, part of our charter is to make what's going to be available. But we're happy to use what's available when it's appropriate, such as bringing population mobility inside NISAC via a technology developed by TRANSIMS. We wonder what this technology can provide us about the location of demographic groups and contact structures—things like that."

Sandia's Snyder says the two labs' modeling approaches are complementary.

"Here, our models are concentrated on infrastructure interdependencies with economics as the base denominator, which looks at how you can determine what the impact of a given scenario is. We're getting a variety of capabilities to bring about a greater whole," Snyder says.

Titan Systems' Smith says simulation designers must be flexible when thinking about possible architectures.

"TRANSIMS is a good example of a small infrastructure solution," he says. "I'm concerned that if you take the solution to a small problem and try to apply it to something so big, you'll end up with something unwieldy. I don't think it's the answer, but I don't think being unfamiliar with it is the answer, either."

Any model must take into account factors beyond physical details, Smith says. Such factors that must be included are the mindset of a given infrastructure's customers and any stockpiles they may have on hand to counteract an outage (see Figure 1).

Misplaced precision?

Another universally agreed-on part of the nation's critical infrastructure is the response of the health care system to a bioterror attack, such as the release of anthrax or smallpox. One longtime modeler of infectious disease, Yale University professor Edward Kaplan, says more emphasis should be placed on broader possibilities for policymakers and less on activity details for theoretical individuals as used in TRANSIMS.

"We have a strong difference of view on where the effort should go," Kaplan says. "My stock phrase is *misplaced precision*. The idea is to get insight. It's not the same model as trying to land a rocket ship on the moon. Decision makers have a limited number of options in front of them."

One of the potential models NISAC officials have spoken of is biosurveillance; an example might be sensors that provide real-time data about medication sales that yield information about a possible epidemic before the medical community can connect the dots through observation and information sharing.

Yet Kaplan contends such an approach is technological overkill. He devised a model that demonstrated that the precision-laden approach originally favored by federal policymakers—where infected individuals' direct contacts would be found and inoculated—would be far less effective in halting an epidemic than a mass vaccination program.

"Our idea is not to replicate precisely how people are interacting with each other in the population," Kaplan says. "Our idea is to model how people are interacting in such a way that the disease transmission that comes out of the model is a conservative representation of what would actually happen. We want to be able to control a worst-case epidemic. It turns out in some ways that the simplest model does the trick."

Kaplan cites material published by British researchers in the journal *Nature*. Models showed that in a "free mixing" model, where each individual in a population interacts randomly with others, each person infected with smallpox would infect between 3.5 and six others during the course of their contagiousness early in an epidemic. In a more structured model, in which people interact in more formal structures (with schoolmates or coworkers, for example, and rarely with people outside their regular areas of activity), the rate of infection is lower.

However, recent results in epidemic modeling have pointed to the "small world phenomenon," Kaplan says, where relatively few interactions between people in different groups can lead to an epidemic almost as severe as a free-mixing model.

"When you look at all this complicated stuff, the question you have to ask is not whether you need to be precise in terms of how people move around, but, rather, are there enough interactions among the groups in the population such that the resulting epidemic looks like one that would come out of this mass action or free-mixing model?" Kaplan says.

He also says that in most bioterror situations, he recommends simulation designers use a simpler-is-better approach.

"I think there are some people at Los Alamos and Sandia who realize they can get more out of the classical models than the fancy ones, that almost every insight they get out of fancy agent-based models, they could have gotten with a simpler model," he says.

Barrett says such arguments miss the point.

"There are many reasons you do this," he says. "Sometimes it's the unavailability of data. You create the data or you otherwise synthesize data that is not available. For example, if you want to know the impact of taking out traffic lights on what mobile communications functions are affected, there's no data."

EDITORIAL OFFICE

COMPUTING in SCIENCE & ENGINEERING
10662 Los Vaqueros Circle, PO Box 3014
Los Alamitos, CA 90720-1314
phone +1 714 821 8380; fax +1 714 821 4010; computer.org/CiSE

Useful URLs

Roger Smith's first counterterror simulation paper:
www.modelbenders.com/papers/siw_s2002

TRANSIMS home page (detailed information on the transportation system simulation model):
http://transims.tsasa.lanl.gov/

Sandia's gateway to NISAC:
www.sandia.gov/CIS/NISAC.htm

He also emphasizes the niche NISAC is expected to occupy shouldn't be painted with an overly broad brush.

"What we produce is analytical capability, sometimes not even study results," he says. "Sometimes we produce tools that others produce study results from. How to triage, how to give up one city for another, is hardly something a computer is going to fix, but we are part of a process of providing more and different information. Pushing a button and having a scientific policy shoot out the other end is not the right model. We're not claiming that, and if anybody thinks that, they should be dissuaded—that's not what we're for."

Greg Goth is a frequent contributor to *CiSE*, *IEEE Software*, and *IEEE Intelligent Systems*.

INTENSITY-MODULATED RADIATION THERAPY

By David I. Lewin

Radiation therapy is a tricky business. The goal is to deliver as much radiation as possible to a tumor while sparing nearby noncancerous tissue—especially critical, but radiation-sensitive organs such as the spinal cord or rectum. To do this, radiation physicists may have to calculate and deliver a concave (kidney-shaped) region of high intensity radiation that wraps around organs at risk. Advances in computing and computer-controlled devices have been critical to the development of this technology.

"The older paradigm of radiation therapy [3D conformal radiation therapy] tried to get a uniform dose of radiation across the tumor and its suspected margins [adjacent noncancerous tissue] with an additional margin to allow for patient motion," says James Deye (pronounced "dye"), chief of the Radiation Research Program in the National Cancer Institute's Division of Cancer Treatment and Diagnosis, in Rockville, Maryland.

A newer approach, intensity-modulated radiation therapy (IMRT), "relaxes the requirement for a uniform dose distribution," he says, yet delivers the radiation more accurately to the cancerous region. This dose distribution, created by small beamlets that can differ in intensity, is expected to produce a treatment field that better conforms to the shape of the tumor.

IMRT and conformal radiation therapy differences

IMRT adds up modulated beams to get a field with variations in intensity in it. Efforts to produce IMRT began in the early 1990s, based on an idea promulgated by Anders Brahme of the Karolinska Institute in Stockholm. Radiation physicist Joseph Deasy of Washington University in St. Louis says that in 1988 Brahme realized that uniform beams could not treat the butterfly-shaped lymph node region in cervical cancer, while beams with nonuniform intensity could. Researchers at Memorial Sloan-Kettering Cancer Center in New York City, M.D. Anderson Cancer Center in Houston, the University of Wisconsin, Madison, and the Nomos Corporation in Sewickley, Pennsylvania, have pioneered the development of IMRT treatment.

In contrast, conformal radiation therapy has its roots in the integration of computed tomography scanners into radiation therapy planning (see "Totally Rad: Planning Radiation Therapy in 3D," Jan./Feb. 1999 issue of *CiSE*). This practice, which goes back to the late 1970s, provides a geometric and density model of the patient (including the tumor), says Arthur Boyer, chair of Radiation Physics at Stanford University Medical Center in Palo Alto, Calif. "Conformal therapy tries to find a set of treatment beams ... without modifying in any way the intensity within the boundaries of the collimation."

Importance of IMRT

Intensity modulation is important for treating brain tumors, which might wrap around the radiation-sensitive spinal cord, or for treating prostate cancer, which wraps around the rectum, Boyer notes. Use of IMRT shields the rectum by compensating for the lack of beams that would have to pass through the rectum to reach the prostate by changing the intensity of beams coming in from the side.

To achieve nonuniform intensity, the beam passes through a dynamic field shaper, such as a multileaf collimator. "Essentially, what one has is a jaw with many leaves" that can move together or apart under computer control, Deasy says. Multileaf collimators can be static (having a fixed configuration during each beam exposure) or dynamic (changing configuration during a beam exposure).

Advances in computing hardware have helped make IMRT possible. The linear accelerators that deliver radiation to cancer patients now rely on embedded computers for control, says Bruce Curran, vice president for technology at Nomos. The ability to dynamically form and modulate radiation beamlets requires the use of real-time feedback loops, Deye notes. Successful use of embedded computers to control such loops requires the use of radiation-hardened components. Nomos, which focuses its product line on IMRT, was "the first company to put out a commercial product," Deye states.

"The bigger jump was the development of inverse planning," Curran says. Planning systems have to be able to manage hundreds of virtual pencil beams of radiation, a task that takes hundreds of thousands of calculations through such Monte Carlo techniques as simulated annealing.

The mathematics of treatment planning is an evolving field. "Some of the mathematics for inverse planning was already at hand, but some was not," Deye says. "We're continuing in our physics research investigating how to optimally plan and deliver IMRT treatments," says Boyer about the Stanford IMRT group.

In February 2002, the National Institutes of Health and the National Science Foundation cosponsored the Operations Research Applications in Radiation Therapy Workshop. This meeting brought operations research experts together with radiation therapists to "have them compare notes about what it means to optimize a treatment plan," Deye says. Right now, op-

BRIEFS

OLD SOFTWARE PROGRAMS LEARN NEW TRICKS

By Anne Jacobson

Like many ambitious workers in the information age, powerhouse software packages learn new skills to keep their resumes fresh in case a better job offer comes along. According to recent trends, good software seems to get around.

A software package designed to process astronomical data from NASA's Hubble Space Telescope recently changed careers. It has gone to work for the Celera Genomics Group and has turned its analytical abilities to Celera's genomic and proteomic data.

Called Operational Pipeline Unified Systems (Opus), the software was first developed by the Space Telescope Science Institute to guide data from Hubble through a processing pipeline that converted the raw information into a form that astronomers could use. Because Celera's large bioinformatics databases are similar to expansive catalogs of astronomical informa-

(Continued on p. 10)

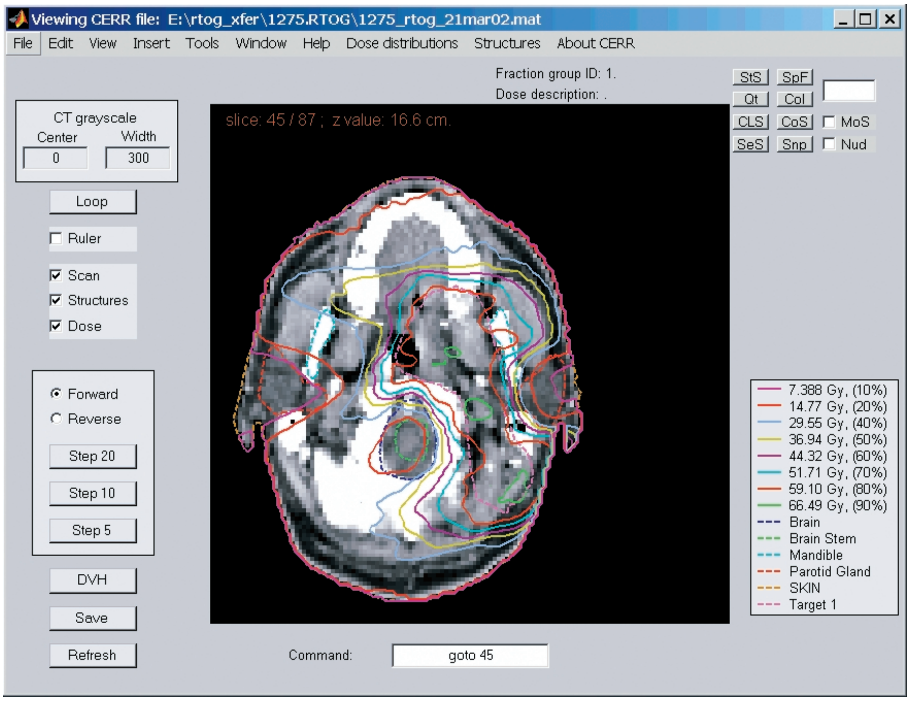

Fraction group ID: 1.
Dose description: .
CT grayscale
Center: 0
Width: 300
slice: 45 / 87 ; z value: 16.6 cm.
Command: goto 45

Figure 1. An image showing regions of equal radiation dose during the treatment of a head-and-neck cancer. Though the use of intensity-modulated radiation therapy (IMRT), the region of the spinal cord receives a lower dose of radiation than does the tumor, even though the tumor wraps partly around the spinal cord. Computational methods and computer-controlled beam collimators calculate and deliver the radiation dose as a series of variable-intensity radiation beamlets. Image courtesy of Joseph Deasy, Washington University at St. Louis Medical Center.

tion, the company licensed Opus in late 2001 to help process and store its data.

Although Opus needed some new on-the-job training—switching from exploring deep space to analyzing the body's tiny secrets—the transition paid off for its new employer. "[Opus] has cut software system deployment time and saved in-house development effort, thus freeing software engineering resources for other tasks," says John Reynders, Celera's vice president for information systems.

Such career switches are common for talented software, says Arnold Peskin, an information technology specialist at the US Department of Energy's Brookhaven National Laboratory. "This is very much the way of computing," he says. "Dual use of software is all over the place."

In another example, the software originally designed to hunt for disease is now searching for oil. As part of the National Gas and Oil Technology Partnership, scientists at Brookhaven National Laboratories team with gas and oil companies to prospect for oil using x rays. First, they shoot x rays into the ground, then they analyze the pattern of the x rays as they bounce back to determine what types of hydrocarbons are present underground and the structure of intervening rock, Peskin explains.

The DOE project uses an algorithm based on the well-known Snark algorithm used in medical imaging technology. Snark was originally developed by researchers at the University of Pennsylvania Department of Radiology to take x-ray data from computerized axial tomography scans and convert that data into visual images of the human body. Brookhaven scientists have adapted the Snark algorithm for their x-ray pattern analysis.

Once a software program proves itself in one field, it is quickly sought after by industry headhunters. "The clear trend is that when a good product becomes available, it gets generalized quickly," Peskin says.

SUPERCOMPUTERS PREDICT SHRINKING BIODIVERSITY IN MEXICO

By Anne Jacobson

Using powerful new software, one of the world's leading databases of species habitats, and an unprecedented collaboration of effort, re-

timization is far from automated, and requires quite a bit of human intervention—"a person has to sit there and assess these plans as they are developed," he says. Radiation physicists need to know the medical aspects of the problems, not just the mathematics and physics, to successfully make a contribution to the field, Boyer notes.

IMRT questions

Although IMRT has gained attention in the radiation oncology world, several important implementation questions remain unresolved, Boyer says. Because the technology is more precise in delivering radiation, there is the need to assure that the target hasn't moved from the time the plan was developed. For lung tumors, this especially includes efforts to compensate for the effect of the patient's respiration. Boyer agrees that this latter effort is, in principle, comparable to the adaptive optics techniques astronomers have developed to compensate for the effect of atmospheric disturbance.

Researchers have expressed a great deal of interest in integrating magnetic resonance imaging and ultrasound to determine the geometric bounds of the tumor, Boyer says. Future investigations include the development of positron emission tomography agents specific for tumor cell surface molecules, or techniques to visualize blood vessel formation or blood flow in tumors.

Another problem is finding the optimal intensity pattern for treating a particular tumor. "This whole issue ... is unresolved," Boyer says. Improved computers and graphics technology are allowing better investigation of the search space, he notes. A great boon has been the adoption of data standards (DICOM-3 for 3D medical imaging, and the evolving DICOM-RT for radiation therapy), which allow different imaging and radiation therapy planning systems to communicate. Continuing work on standards "would be quite helpful," he says.

Where is IMRT going? The recent decision by the US Centers for Medicare and Medicaid to allow payment for IMRT has given the technique legitimacy, Deye says. "The technology's status is definitely in a state of flux," Boyer says. "The question is, how fast is it being assimilated by the community hospitals and freestanding radiation therapy facilities?"

David I. Lewin is contributing editor to Computing in Science & Engineering, IEEE Intelligent Systems, and Distributed Systems Online. Trained in biology, he writes on science, technology, and medicine from Silver Spring, MD. Contact him at lewinwriter@earthlink.net.

searchers have predicted that greenhouses gas and other ecosystem threats will bring great instability to Mexican wildlife.

"This research marks a major step forward in being able to investigate in a quantitative way the initial impacts of climate change on ecosystems and biodiversity," says lead researcher A. Townsend Peterson of the University of Kansas Natural History Museum and Biodiversity Research Center.

"What's unique about what we've done is that for the first time we were able to look at a whole community in detail," he says. The analysis is the largest of its kind, covering the entire country of Mexico and examining over 1,800 species, including all of the country's mammals and birds and many of its butterflies.

Key to the effort was a software program called the Genetic Algorithm for Rule-Set Prediction (Garp), created by David Stockwell of the San Diego Supercomputer Center at the University of California, San Diego. Garp uses machine-learning methods from artificial intelligence research, called genetic algorithms, to generate and test scenarios about how species would migrate in response to higher average temperatures and less rainfall.

Over the next 50 years, global warming will force the majority of species in Mexico into narrower geographical ranges, Garp predicts. In addition, climate change will throw new predators and prey together and allow new diseases and parasites to spread.

The analysis took about 30 days, running on a supercomputer provided by the National Partnership for Advanced Computational Infrastructure. Results of the research, supported by the National Science Foundation and the Mexican government, appeared in a recent issue of *Nature*.

Garp used species habitat information from a database assembled by Mexico's biodiversity commission, Conabio. The database contains information from more than 40 natural history museums around the world.

"We used more than 110,000 unique records of observed species locations," says Peterson. "That species-by-species look has allowed us to appreciate just how big the differences are in the way each species can respond to climate change—it's quite complex and individual."

Next, the researchers plan to produce computer models that explore the combined effects of climate change and other scenarios such as invasive species on biodiversity.