

## Modeling Chaotic Storms

*Scientists say improvements to extreme-weather prediction are possible with new weather models and a reinvention of the modeling technologies used to process them.*

**T**HE WARNING TIME for the onset of extreme weather is far greater today than it was 20 years ago, when only a few minutes of warning could be given for tornadoes, and only half of them could even be predicted. Today, new data-collection technologies, such as Doppler radar and satellites, have improved the ability to identify and track hazardous weather. But scientists say further improvements to warnings' lead times will not come primarily from the physical systems that gather weather data, but from improving the modeling technologies used to process the data, and from improving the prediction models themselves.

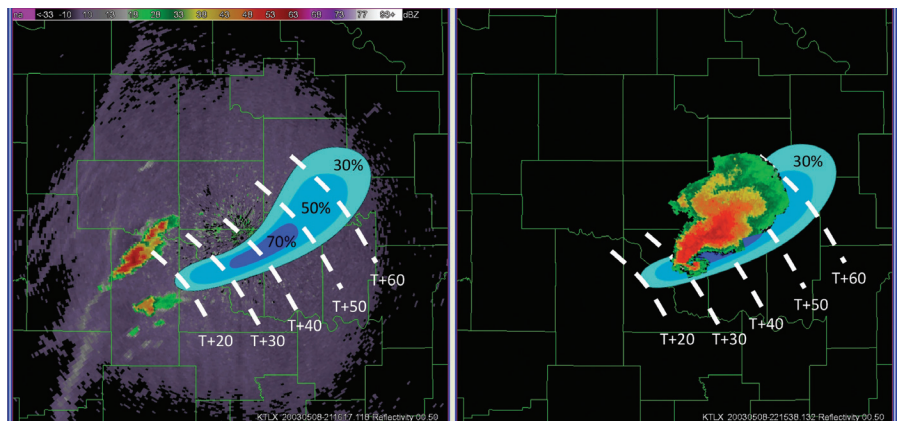
Even with the best data-collection technology, weather-tracking systems cannot be effective at predicting what might come next without refined prediction models that take into account advanced physics and other factors that make running such models a task that takes so long, even on the fastest supercomputers, that the results are not timely enough to be useful. Steve Koch, who was the director of the Global Systems Division at the U.S. National Oceanic and Atmospheric Ad-

ministration's (NOAA's) Earth Systems Research Laboratory (ESRL) before becoming director of its National Severe Storms Lab (NSSL), is focused on developing advanced computing architectures to improve this situation.

"It looks like we are approaching an average lead time of about 14 minutes of warning for tornadoes, and we are reaching the limit of how far we can push detection with technology," says Koch. "We might get perhaps 20

minutes by using more advanced radar technologies, but ultimately if we want to get to a one-hour forecast of a tornado, we cannot do that right now just with radar."

Koch, who has been working at NOAA for more than a decade, says the organization is focusing extensively on global models, not only of weather but also of climate and the influence of oceans, land surfaces, ice and snow, aerosols, and other factors, including everything from anthropogenic effects on climate to invasive species in the Great Lakes. As a result, this focus has produced what Koch calls an "explosion of modeling" in the past few years. "There is great interest in integrating as many of these processes as possible into a coherent, scientific approach to the problem of forecasting the Earth system," he says.



**In the Warn-on-Forecast system being developed at the National Severe Storms Laboratory, an emerging thunderstorm is observed by radar (left). That radar data is used as input for an ensemble of prediction models to determine the probability of a tornado during the next hour. The blue shading represents tornado probability, while the white dashed lines indicate the storm location in minutes from the present. Thirty minutes later (right), a thunderstorm is observed by radar over the area of greatest tornado probability, as predicted.**

However, despite the proliferation of such models, the idea of a completely integrated approach, says Koch, is something that NOAA is still working out. “We don’t have one gigantic model that solves all the problems,” he says. Even so, one result of this intensive focus on modeling is that NOAA has made measurable progress at improving the ability to predict extreme weather more accurately. However, at NSSL, the focus of which is tornadoes, hail, and high winds, the difficulty of such predictions is acute.

Until recently, weather prediction has mostly relied on a few central prediction models, but that strategy has given way to running an ensemble of models, sometimes upward of 40 models, each with slightly different physics or other configurations. The idea is to create models that can capture the wide variability in the atmosphere and then combine them so there is sufficient spread in the characteristics of a complete ensemble. The ultimate goal of ensemble modeling, then, is to develop a broader representation of plausible outcomes while also producing useful estimates for forecasting.

One problem, however, is that even with an ensemble of models, the prediction results are limited by the quantity and quality of input data. In the U.S., for example, 142 weather radars cover the lower 48 states, with an average spacing of approximately 260 kilometers between them. Because of the Earth’s curvature, more

**Until recently, weather prediction has largely relied on a few central prediction models, but that strategy has changed to running an ensemble of models, each with slightly different physics or other configurations.**

than 70% of the U.S. is not covered by radar in the lowest 1 kilometer of the atmosphere, which is, of course, where humans live. It is also where severe thunderstorms develop. “To overcome the Earth’s curvature problem,” says Koch, “you would need to develop a cost-effective, much-denser coverage system.”

Another challenge is the weather processes. What creates severe turbulence, for example, is far from being completely understood. “There are gaps in our understanding of processes within thunderstorms, and one of the biggest gaps is the physics of

the precipitation process itself,” says Koch. “We have an unsatisfactory understanding of what generates organized thunderstorms.”

#### **A Matter of Volume**

Beyond these issues, which have drawn much attention by researchers in recent years, the sheer volume of data produced by an ensemble model configuration, which scientists now believe is the most effective way to predict extreme weather, is much more than the bandwidth of the most advanced networks can handle. In addition, while an ensemble of models can be tuned to the particular aspect of the weather it is designed to simulate, rendering the ensemble quickly enough to be useful for real-time forecasting remains a challenge, even with today’s fastest supercomputers.

The power and cooling requirements are formidable for supercomputer-class systems, and the maintenance costs can be exorbitant. So Koch and other NOAA scientists are looking beyond traditional supercomputer systems and are exploring systems based on graphical processing units (GPUs). While GPUs are not general purpose, they do offer performance advantages over general-purpose CPUs for certain applications. NOAA has not made a decision yet about implementing GPU systems, but trials are in an advanced exploratory phase run at NOAA’s ESRL.

Mark Govett, chief of the advanced computing lab at ESRL, has been di-

## **Technology**

# **PC and Tablet Sales Forecast**

Industry analysts are less optimistic now than earlier this year about PC sales for 2011. But the personal computer’s woes have little to do with the public’s enthusiasm for tablets as some believe, says Stephen Baker, vice president of industry analysis for The NPD Group.

Research company Gartner has trimmed its forecast for 2011 PC sales from 9.8% growth just three months ago to 3.8%, while rival research firm IDC lowered its forecast from 7.1% to 4.2%.

The NPD Group would not release specific numbers, but Baker characterizes U.S. PC sales as having been mostly flat to slightly down, mainly due to the weak economy. “And also because people bought record numbers of computers these last two years ever since the launch of Windows 7,” Baker says. “Those new computers won’t need to be replaced or upgraded for a while.”

On the other hand, tablet sales—which have been

dominated by Apple’s iPads—have grown dramatically from a first-year small base, Baker says. There was no tablet market until the iPad was released in April 2010.

But, Baker adds, \$600 iPads have had very little effect on the sales of less-pricey PCs that average under \$500 for notebooks and about \$500 for desktops. “Also, people typically buy tablets for different reasons than for PCs, mostly for content consumption, like watching

movies, listening to music, and surfing the Net, while they buy PCs for content creation and productivity, such as editing pictures, creating movies, and managing finances.”

Baker believes these trends will continue until at least 2012’s second half when Microsoft is expected to unveil Windows 8. “That’s when we foresee some changes in PC momentum,” he says, “which will start to put some pressure on the iPad.”


—Paul Hyman

recting these GPU trials. “We’re always looking for cost-effective computing that can run research weather models,” says Govett, who describes GPUs as the next generation of supercomputing technology. Govett says that, to run the new weather models, NOAA would need traditional CPU-based systems that would cost \$75 to \$150 million and would require building special data facilities to accommodate the power and cooling. “GPU systems with similar capabilities could be built for about \$10 million, with no special building needed,” he says.

As an example of the kind of cost savings involved with GPUs, Govett cites a test run with a next-generation model called the nonhydrostatic icosahedral model (NIM). The advanced computing lab at ESRL, which includes model developers, code-parallelization experts, and GPU researchers, calculates that more than 200,000 CPU cores would be needed to produce a NIM forecast close enough to real-time to be useful for prediction. The ESRL researchers, who began experimenting with GPUs in 2008, demonstrated that the NIM model could be run 25 times more quickly on GPUs than on traditional CPUs. (For an overview of new developments in high-end computing with GPUs, see “Supercomputing’s Exaflop Target” in the August 2011 issue of *Communications*.)

While GPUs appear to be a promising alternative to traditional supercomputing, several challenges could prevent them from being adopted for weather modeling. For one, the code must be modified to run on GPUs. Govett and his team have developed their own compilers to convert their Fortran code into the language used on NVIDIA GPUs. As these compilers mature, Govett explains, the parallelization process will get easier. But, at least for now, Govett calls the work to parallelize models to run efficiently on GPUs a significant challenge. That code-parallelization difficulty is magnified with the new class of ensemble modeling that is expected to revolutionize weather prediction.

Single weather models, by themselves, present a significant challenge to high-performance systems capable of handling extreme workloads. Large ensembles consisting of many mem-



**A hexagonal grid used by several next-generation global weather models. This particular grid is based on a 480-kilometer model. The next generation of weather models, driven by GPU technology, will be run at a scale of 2 to 4 kilometers, making neighborhood-level accuracy possible.**

bers (or models with different configurations) generate a range of forecasts that are then combined to produce a single more accurate forecast. At the finest scale needed for accurate weather prediction, these ensembles can be run only on the fastest supercomputers. “We recently ran a 20-member ensemble on the Oak Ridge Jaguar supercomputer, which was the largest supercomputer in the world until last year,” says Govett. “That model required over 120,000 CPU cores, or basically half of the machine, and this was for one ensemble run.”

Govett says NOAA does not have the processing power to run large ensemble models at a research level, let alone at an operational level where the models are run quickly enough for the predictions to be useful for forecasting. “Research computing, and to some extent climate forecasting, does not have the time constraint that operational weather forecasting does,” says Govett. “In operations, weather models need to be run quickly or the information they produce will not be useful, particularly for severe weather where lives and property are at risk.”

As for the future of ensemble modeling, Govett says that, by exploiting the parallelism available in GPU and multicore systems and rewriting existing code with new algorithms and solvers, ensemble models will be able to run at a global scale and generate weather and climate predictions with much better accuracy than can be

achieved today. Despite the ongoing challenges facing Govett and other researchers working to refine weather modeling and find alternatives to traditional supercomputer-class systems so the more sophisticated ensemble models can move from research to operational use, Govett says he remains optimistic about the years ahead. “Model prediction continues to improve,” he says. “This is particularly evident in the improved accuracy of hurricane track and intensity forecasts, severe weather predictions of flooding and tornadoes, and regional climate prediction.”

Koch, for his part, says the future of weather prediction looks bright, indeed. But getting to the point where scientists are able to produce 60-minute warnings for extreme weather, he says, will be a major undertaking that will require enough computing power to run fine-scale ensemble models at an operational level. “That’s my dream,” he says. “It may be achievable within 15 years.”

#### Further Reading

Govett, M., Middlecoff, J., and Henderson, T. Running the NIM next-generation weather model on GPUs, *Proceedings of the IEEE/ACM International Conference on Cluster, Cloud, and Grid Computing*, Melbourne, Victoria, Australia, May 17–20, 2010.

Henderson, T., Govett, M., Middlecoff, J., Madden, P., and Rosinski, J. Experiences applying Fortran GPU compilers to numerical weather prediction models, *Proceedings of the 2010 Symposium on Application Accelerators in High Performance Computing*, Knoxville, TN, July 13–15, 2010.

Stensrud, D.J., et al. Convective-scale warn-on-forecast: A vision for 2020, *Bulletin of the American Meteorological Society* 90, 10, Oct. 2009.

Stensrud, D.J. and Gao, J. Importance of horizontally inhomogeneous environmental initial conditions to ensemble storm-scale radar data assimilation and very short range forecasts, *Monthly Weather Review* 138, 4, April 2010.

Yussouf, N. and Stensrud, D.J. Impact of high temporal frequency phased array radar data to storm-scale ensemble data assimilation using observation system simulation experiments, *Monthly Weather Review* 138, 2, Feb. 2010.

**Kirk L. Kroeker** works in communications and has written extensively about the impact of emerging technologies.

© 2011 ACM 0001-0782/11/11 \$10.00