

SETTING THE STAGE

In 1989, when CAPS started its work, the prevailing view about storm-scale weather forecasting was skepticism; thunderstorms were thought to be inherently chaotic and unpredictable. Much has changed. CAPS developed innovative techniques to gather atmospheric data from Doppler radar, and they developed a forecast model to use this data at the scale of thunderstorms. In the mid-90s, via a series of spring collaborations with PSC, CAPS proved the feasibility of forecasting thunderstorms.

In 2005, in a large-scale collaborative experiment with NOAA, again with major support from PSC, CAPS took another leap forward by showing that it was possible, in some circumstances, to predict the details of a thunderstorm as much as a full day in advance. "We provided dramatic new evidence that the predictability of organized deep convection is, in some cases, an order of magnitude longer — up to 24 hours — than suggested by prevailing theories of atmospheric predictability," says former CAPS director Kelvin Droegemeier.

An opportunity to build on this prior work, the 2007 HWT Experiment involved more than 60 researchers and forecasters from government agencies, universities and the private sector. Along with PSC, CAPS and SPC, other collaborators were the NOAA National Severe Storms Laboratory in Norman; the NOAA National Centers for Environmental Prediction Environmental Modeling Center; the National Center for Atmospheric Research; LEAD (Linked Environments for Atmospheric Discovery), an NSF Large Information Technology Research grant program and TeraGrid Science Gateway; and the National Center for Supercomputing Applications (NCSA) in Illinois, a lead TeraGrid resource provider.

To implement the CAPS daily forecast runs using the WRF (Weather Research and Forecast) model on the XT3, PSC provided technological and staff assistance at several levels:

- PSC networking staff coordinated with OneNet, a regional network of the State of Oklahoma, and National Lambda Rail (NLR), a network initiative of U.S. universities, and with Cisco Systems, who contributed a dedicated "lambda" [a 10-gigabit-per-second optical-network] for up to a 12-month period.
- PSC implemented and began testing the lambda at its end in January, using existing equipment in the Pittsburgh metro and local-area network. The backbone was provided by NLR and OneNet provided the link from Tulsa to Norman, Oklahoma.
- This dedicated link — from the Cray XT3 to OneNet in Tulsa to a supercomputer at the University of Oklahoma (which ingested and post-processed the data) — made possible the transfer of 2.6 terabytes of data per forecast day.
- PSC staff optimized the latest version of the WRF model to run on the Cray XT3, gaining a threefold speedup in input/output (I/O) of the WRF code, substantially improving overall performance.
- PSC also optimized the I/O for post-processing routines used to visualize and analyze the forecast output, achieving 100-fold speedup.
- PSC modified the reservation and job-processing logic of its job-scheduling software to automatically schedule the WRF runs and related post-processing, 760 separate jobs each day, demonstrating the TeraGrid's ability to use the Cray XT3, a very large "capability" resource, on a scheduled, real-time basis.

ANOTHER LEAP FORWARD

The daily runs of the 10-member WRF ensemble each used 66 processors of PSC's Cray XT3 for 6.5 to 9.5 hours each day. The single high-resolution WRF forecast used 600 XT3 processors for about nine hours daily, including data dumps at five-minute intervals via the dedicated network to a system in Norman, where CAPS post-processed the data for visualization and analysis.

It was an unprecedented marshaling of computing and network capability. "To do this amount of processing with the ability to transfer these huge volumes of data from Pittsburgh to Norman on a reliable, daily, timely basis was a technological success that we're still appreciating," says Weiss. "It hasn't been done before. There were people who thought 'How are they going to do that?'"

"IT HASN'T BEEN DONE BEFORE. THERE WERE PEOPLE WHO THOUGHT 'HOW ARE THEY GOING TO DO THAT?'"

Over its eight weeks, the HWT experiment produced massive amounts of data, and the analysis is ongoing. "Preliminary findings," says Xue, "show successful predictions of the overall pattern and evolution of many of the convective-scale features, sometimes out to the second day. The ensemble shows good ability in capturing storm-scale uncertainties."

The case of a large-scale north-south storm-front that developed on May 23 and continued for several days across the Great Plains exemplifies the ability of the storm-scale runs to produce a forecast that holds up for the storm system as a whole as much as 33 hours in advance. "The rather successful prediction of the overall pattern and evolution of this squall-line case," says Xue, "is very encouraging, especially considering that we are examining convective-scale predictions for up to 33 hours."

The HWT experiment also used capabilities developed by LEAD (Linked Environments for Atmospheric Discovery), a TeraGrid Science Gateway, to test "on-demand" forecasts. Triggered automatically from daily SPC forecasts indicating regions within the overall HWT domain where storms were likely to develop, these forecasts used NCSA's Ensemble Broker software and its 16-teraflop Tungsten system to run at two-kilometer resolution within the regions of high storm likelihood. More than 1,000 of these on-demand forecasts ran successfully.

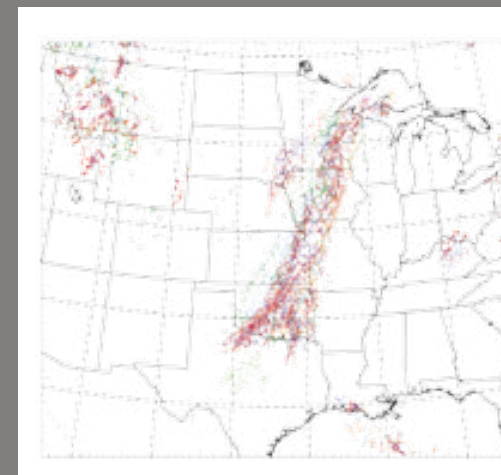
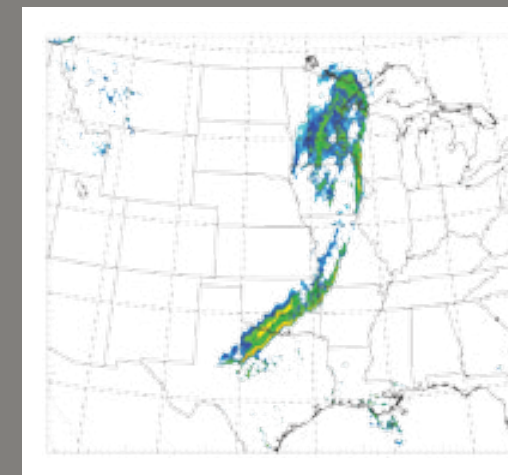

"This experiment was an enormous leap forward," says Droegemeier, who directs LEAD and as former CAPS director has led several other storm-forecast experiments over the past decade.

MORE INFORMATION:

www.psc.edu/science/2007/storms.html

Ensemble Predictions: Spaghetti Plots & Postage Stamps

This plot (left) from the forecast for noon central standard time (18 universal time) on May 24, 2007 (from a model run on May 23) shows probability of radar reflectivity (proportional to intensity of precipitation) derived from the 10-member ensemble forecast 21 hours in advance compared to actual observed radar reflectivity on May 24 at noon CST (right).



The "spaghetti plot" (left) for May 24 at noon CST represents the distribution of the radar-reflectivity forecast among the 10 ensemble runs, each represented by a different color. The corresponding "postage stamp" (below) display shows each of the 10 ensemble forecasts. Differences among them comprise a quantifiable range of the uncertainty associated with model errors and limitations on storm-scale predictability, as opposed to the traditional single-member (deterministic) forecast that produces only one prediction, good or bad.

