


# The future of forecasting: Leading weather agency turns to Titan to advance science of prediction

August 21 2015, by Gregory Scott Jones And Jeremy Rumsey

---



The IFS alternative dynamical core option: Left, an example of an unstructured mesh for a low-resolution model. Right, the domain decomposition used in IFS; each patch represents the grid area owned by an MPI task. Credit: Oak Ridge National Laboratory

Knowing how the weather will behave in the near future is indispensable for countless human endeavors.

According to the National Oceanic and Atmospheric Administration, [extreme weather events](#) have caused more than \$1 trillion in devastation since 1980 in the United States alone. It's a staggering figure, but not nearly as staggering as the death toll associated with these events—approximately 10,000 lives.

The prediction of low-probability, high-impact events such as hurricanes, droughts, and tornadoes, etc., has proven to have profound economic and social impacts when it comes to limiting or preventing mass property damages and saving human lives. But regardless of the aim, predicting weather has always been a tricky business.

However, thanks to one of the world's most powerful computers, it's becoming less tricky and more accurate. Researchers from the European Centre for Medium-Range Weather Forecasts (ECMWF) have used the Titan supercomputer, located at the US Department of Energy's (DOE's) Oak Ridge National Laboratory, to refine their highly lauded weather prediction [model](#), the Integrated Forecasting System (IFS), in hopes of further understanding their future computational needs for more localized weather forecasts.

ECMWF is both a research institute and a 24/7 operational service, supported by 34 European countries. In the United States, the IFS is perhaps best known as the weather model that gave the earliest indication of Hurricane Sandy's path in 2012. Sandy is the second costliest hurricane in US history and the most powerful of the 2012 season. "Our ensemble forecasting system predicted the landfall of superstorm Sandy on the US East Coast more than 7 days in advance," said Erland Kallen, director of research at ECMWF.

Sandy was responsible for 117 American deaths (primarily because of drowning) and cost the United States \$65 billion. Obviously, the sooner decision makers know where storms like Sandy are heading, the faster

they can move to minimize the damage.

And that's precisely the aim of ECMWF's research on Titan at the Oak Ridge Leadership Computing Facility, a DOE Office of Science User Facility. By harnessing 230,000 of Titan's cores and refining communication within their model, researchers have been able to determine the limits of their current model implementation and the computational refactoring necessary to take the model to unprecedented levels of resolution and detail.

The research is part of a broader drive to increase the scalability of the IFS, which produces forecasts every 12 hours and predicts global weather patterns up to 2 weeks ahead. It was carried out under a European Union-funded project known as Collaborative Research into Exascale Systemware, Tools, and Applications, or CRESTA. This project, which finished at the end of 2014, aimed to prime certain applications for future computer architectures, up to and including the exascale. It also involved the supercomputer manufacturer Cray (Titan's manufacturer), Allinea (the developer of the DDT debugger), and a number of European universities and high-performance computing (HPC) centers.

"ECMWF is looking at the scalability of the IFS over the long term," Kallen said. "As well as operational forecasting, we also do cutting-edge research which feeds back into regular upgrades of the IFS."

## **Meteorological methodologies**

Because physical laws govern weather, changes in parameters such as temperature, wind, and humidity can be described with mathematical equations. However, these non-linear partial differential equations are highly complex and provide a range of solutions depending on model uncertainties and uncertainties in the initial state from which they are

advanced in time.

The IFS algorithm incorporates a spectral transform method—a numerical technique that converts data on a model grid into a set of waves. Uncertainties in the model and in the observed atmospheric state are taken into account by forming an ensemble of solutions. This provides precious information on the degree of confidence associated with particular predictions. Over the last 30 years, this widely used forecasting method has proven to be highly scalable and to work well with most HPC architectures.

Using Titan, the team was able to run the IFS model at the highest resolution ever (2.5-kilometer global resolution), using Coarray Fortran, an extension of Fortran, to optimize communication between nodes. This brought a major improvement in performance with only a modest introduction of new code.

The team also achieved the benchmark of simulating a 10-day forecast in under one hour using a 5-kilometer global resolution model on 80,000 Titan cores. Without the addition of coarrays, the simulations would have required up to 120,000 cores. This meant that, by reworking basic communication procedures, the team was able to achieve a 33 percent increase in performance—a vital improvement necessary for meeting the team's future scalability goals.

The ECMWF team now has a "roadmap" to assist it in planning, particularly in terms of scalability on future architectures. The roadmap includes dates on which higher-resolution models are expected to be introduced, said George Mozdzynski, a consultant at ECMWF and team member on the Titan research project. In the past, ECMWF has doubled the resolution of its model roughly every 8 years, and this rate forms the basis of the roadmap. With new architectures featuring accelerators such as Titan's GPUs, this roadmap could possibly be accelerated.

"We have learned a lot, and we now know where the bottlenecks lie as we scale up our model to next-generation systems," Mozdzynski said.

"The spectral transform method we've been using for 30 years is now experiencing some limitations at scale, as a result of the fact that it's heavily dependent on communication. There are very high message passing costs, and that's something that's unlikely to improve over time."

## **The computation/communication combination**

The IFS forecast model relies heavily on the conventional Message Passing Interface (MPI), the mechanism responsible for code communication in countless applications. As models are scaled up in resolution, there needs to be a corresponding reduction in the model time-step, and at some point MPI communications can become crippling and stall performance. In fact, these communications can account for up to 50 percent of the IFS's execution time at high core counts.

Although ECMWF has yet to harness the power of Titan's GPUs, the traditional CPUs were sufficient to identify scaling issues in their model.

To address the communication bottleneck, the team found a solution in using coarrays, which are part of the Fortran 2008 standard and fully supported by the Cray Fortran compiler on Titan.

In a system as large as Titan, all of the nodes are engaged in similar calculations, such as arriving at updated values for pressure and temperature. To do this, they all need information from the neighboring nodes, which means there is constant communication between the nodes themselves and within individual nodes.

Applications such as ECMWF's IFS model use OpenMP (a standard for shared memory parallelization) for computations on the nodes and MPI for communications across the nodes. Generally, this arrangement works

well for many applications, but for high-resolution weather modeling using the spectral transform method, it can bog down the application, wasting precious runtime—it simply does not easily allow for any overlap of computation and communication. The use of coarrays (also called PGAS, for Partitioned Global Address Space) in Fortran, however, allows a simple syntax for data to be communicated between tasks at the same time computations are being performed.

"We have overlapped the communication with computation instead of having just the OpenMP parallel region for the computation," Mozdzyński said. "Not everything was perfect when we started using coarrays, but over time each compiler release got better, so much so that at this time everything is working perfectly with our coarray implementation."

Essentially, by overlapping computation and communication, the team's processing times have decreased dramatically. While this is only necessary at scale and only with a 5-kilometer model, or into the 2020s, the team already has begun to solve the processing time problem.

"We believe that coarrays will become increasingly used at core counts where MPI no longer scales," Mozdzyński said. "While MPI is a very good standard, and it will continue to be used for decades to come, in time-critical code sections there can be other approaches, such as PGAS, that we believe will become increasingly used."

## **Beyond the near future**

While this new approach bodes well for the near future, other changes will be necessary as ECMWF continues to push the envelope in weather forecasting.

For instance, the team currently can run a 2.5-kilometer model on Titan



but can only accomplish 90 forecast-days per day. Ideally, they would like to get to 240, or the equivalent of a 10-day forecast in one wall-clock hour. "With the current method, our performance requirements are not being met. We need more flexibility to use accelerators and new algorithms than the current model gives us," said Mozdzynski, adding that development of an alternative "dynamical core" option in IFS—to include algorithmic and grid choices used to integrate dynamical gas-flow equations describing mass conversion, momentum, and energy, in both vertical and horizontal directions, on the surface of the rotating globe—is 5 or more years down the road. "The future looks good, but there's lots of work to do."

Ultimately, new tools and ways of thinking will be necessary to move far into the future. Mozdzynski points out that the use of unstructured meshes in the spherical model with new numerics, along with a number of vendor tools, are all being explored to move the science forward.

Before Titan, the team had only run the IFS forecast model on 10,000 cores and was unsure of how the model would react to such a massive system. "The experience that we've gained running on an extreme number of cores is really remarkable," Mozdzynski said. "We were almost certain that at some point it would blow up, but it all came together, which is rather surprising."

The success of the model modified to work with the new communication framework has given the team confidence to move forward into the realm of the world's most powerful computers. With sufficient computing resources, Mozdzynski believes the team could run the 5-kilometer model today if necessary. "We've already proven we can do it," he said.

Although more work needs to be done to prepare numerical weather prediction for future computer architectures, the team has already made

excellent progress—partly thanks to Titan.

Provided by Oak Ridge National Laboratory

Citation: The future of forecasting: Leading weather agency turns to Titan to advance science of prediction (2015, August 21) retrieved 7 May 2024 from <https://phys.org/news/2015-08-future-weather-agency-titan-advance.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.