CASE STUDY

Optimizations of the Unified Model Code on the Cray® XC30™ Supercomputer "ARCHER" Yield Up to 16 Percent Speedup

Organizations
U.K. Met Office
www.metoffice.gov.uk

National Centre for Atmospheric Science
Computational Modeling Services
www.ncas.ac.uk

Segment
Weather

Application
Optimization testing of the Met Office Unified Model (UM) on a Cray XC30 supercomputer

"CrayPAT™ and Cray Reveal have proved invaluable as tools to help us both diagnose sub-optimal model performance and efficiently apply directives to improve thread load balancing. Squeezing every ounce of performance out of our computationally very expensive integrations means more science per ARCHER core-hour — good news all round."

—Dr. Grenville Lister,
Head of Computational Modelling Support,
National Centre for Atmospheric Science (NCAS)

Test Results
Using a Cray performance tool named Reveal to add OpenMP directives to UM, researchers increased speed by up to 16 percent on the U.K. National Supercomputer Service's Cray XC30 system "ARCHER" while it slowed performance on the IBM Power 755 MONSooN system. These performance optimizations have resulted in savings of tens of millions of core hours on current climate projects.

Cray XC30 "ARCHER" Supercomputer
ARCHER is a Cray XC30 supercomputer, deployed at the Edinburgh Parallel Computing Centre (EPCC) as the U.K. national computing facility for environmental and engineering science. The 26-cabinet, 2.5-petaflop system provides a capability resource to allow researchers to run simulations and calculations that require large numbers of processing cores working in a tightly coupled, parallel fashion. ARCHER features 9,840 12-core Intel® Xeon® E5 processors and a Cray Aries interconnect.

Cray Inc.
901 Fifth Avenue, Suite 1000
Seattle, WA 98164
Tel: 206.701.2000
Fax: 206.701.2500
www.cray.com

Collaborating to Boost Unified Model Performance
The Met Office Unified Model (UM) is a numerical weather prediction and climate modeling software suite used for simulating weather, climate and earth system processes. Originally developed by the U.K. Met Office as its main forecasting model, the code now has users and development collaborators from weather and climate research centers worldwide.

One such contributor is the U.K. National Centre for Atmospheric Science Computational Modeling Services (NCAS-CMS). NCAS-CMS supports climate investigation by the U.K. academic community, including study of climate change, air quality and hazardous weather, and technologies for observing and modeling the atmosphere. The group conducted UM code optimization work on the U.K. National Supercomputing Service's Cray® XC30™ supercomputer named "ARCHER."

Using Cray Performance Analysis Tools (CrayPAT™) and Cray Reveal to analyze and optimize UM performance, the team was able to automatically design and insert OpenMP directives in the UM, which resulted in speed increases of up to 16 percent.

The performance optimizations on the Cray XC30 system resulted in savings of tens of millions of core-hours on current climate projects.

Optimization Parameters
The primary versions of the shared UM code are currently maintained on an IBM Power 755 supercomputer named MONSooN at the Met Office (to be replaced by a Cray® XC40™ system in late 2015) and codes are typically ported from that environment onto target platforms.

For the NCAS-CMS project, researchers Karthee Sivalingam and Grenville Lister ported UM from the MONSooN environment onto the ARCHER platform, examined its performance at three different resolutions (N96, N216 and N512, corresponding to ≈ 130, 60 and 25 km grid spacing). The objective was to compare performance of UM at a range of scales, and with a range of optimizations, including altered MPI rank placement and the addition of OpenMP directives.

The testing took place in two phases. First, using CrayPAT, the team analyzed the baseline performance of UM to understand the bottlenecks of high-resolution versus low-resolution jobs and how the profile changes as N512 scales up. Then, the team used Cray Reveal to identify and parallelize serial loops using OpenMP directives.

Results

The team's performance analysis using CrayPAT showed that MPI communication and thread imbalance affect the scaling of high-resolution UM jobs. First, message passing overheads increased by 13 percent as the number of nodes increased from 73 to 241. CrayPAT can be configured to detect the MPI grid and MPI communication pattern used and suggest MPI rank placements that can reduce off-node traffic and increase overall performance. These suggestions yielded UM performance improvements by between 5 and 12 percent compared to default SMP rank order.

Next, the team addressed the thread imbalance. OpenMP provides a standard and portable way of parallelizing loops. It involves scoping the loop variable and inserting OpenMP directives before a loop. However, parallelizing all the loops is expensive and requires careful consideration of data consistency, and race conditions involving parallel threads are difficult to debug. With more than 2,000 loops contributing significantly to the runtime of the UM, this undertaking is huge. Cray Reveal, an integrated performance analysis and code optimization tool, aids in this process by providing loop analysis and scoping of serial loops and suggesting OpenMP directives that can be inserted to a loop. For this test, the team parallelized 2,389 serial loops by adding OpenMP directives as suggested by Cray Reveal.

On ARCHER, the UM code achieved a speed up of 16 percent on the N512 job. Overall, the researchers concluded that the investment of months in analysis and optimization yielded performance gains that correspond to the saving of tens of millions of core hours on current climate projects.

ARCHER SURGES AHEAD: Optimization work performed by the National Centre for Atmospheric Science sped up the Unified Model code on Cray XC30 supercomputer ARCHER by 16 percent on a high-resolution model.

About Cray Reveal

Cray Reveal is a performance analysis and code optimization tool. Reveal assists users optimizing code by providing variable scoping feedback and suggested OpenMP compiler directives. Most applications running on the current generation of HPC systems use a distributed programming model like MPI or UPC, but don't take advantage of shared memory programming within a node. To efficiently use resources available on nodes with many cores, application developers must incorporate additional levels of parallelism, such as OpenMP. Adding parallelism to the most important loops within an application is key to making good use of available cores and achieving maximum performance. However, this difficult task involves analyzing code for dependencies and scoping many variables through deep call chains. Reveal extends Cray's existing performance measurement, analysis and visualization technology by combining runtime performance statistics and program source code visualization with Cray compiler optimization feedback. With Reveal, the developer can browse source code, quickly retrieve compiler optimization information, determine which high-level loops can most benefit from OpenMP parallelism, scope variables through complex call chains, and form OpenMP directives.