HYPERION RESEARCH

HPC User Forum Update

Enabling Performance Portable Climate Simulations on Aurora and Frontier

Bob Sorensen, Melissa Riddle
February 2020

IN THIS UPDATE

The HPC User Forum was established in 1999 to promote the health of the global HPC industry and address issues of common concern to users. In September 2019, the 73rd HPC User Forum took place at Argonne National Laboratory in Illinois. This update summarizes a presentation from that meeting entitled *Enabling Performance Portable Climate Simulations on Aurora and Frontier* by Nichols A. Romero from Argonne National Laboratory.

Romero's talk described efforts to port global earth system models, like E3SM, to accelerated pre-exascale DOE systems. Originally written in Fortran, these E3SM codes need to be optimized for GPUs in order to take advantage of the Summit, Frontier, and Perlmutter supercomputer architectures. The development is driven by the DOE-SC mission to understand planetary energy and water issues for the next 40 years.

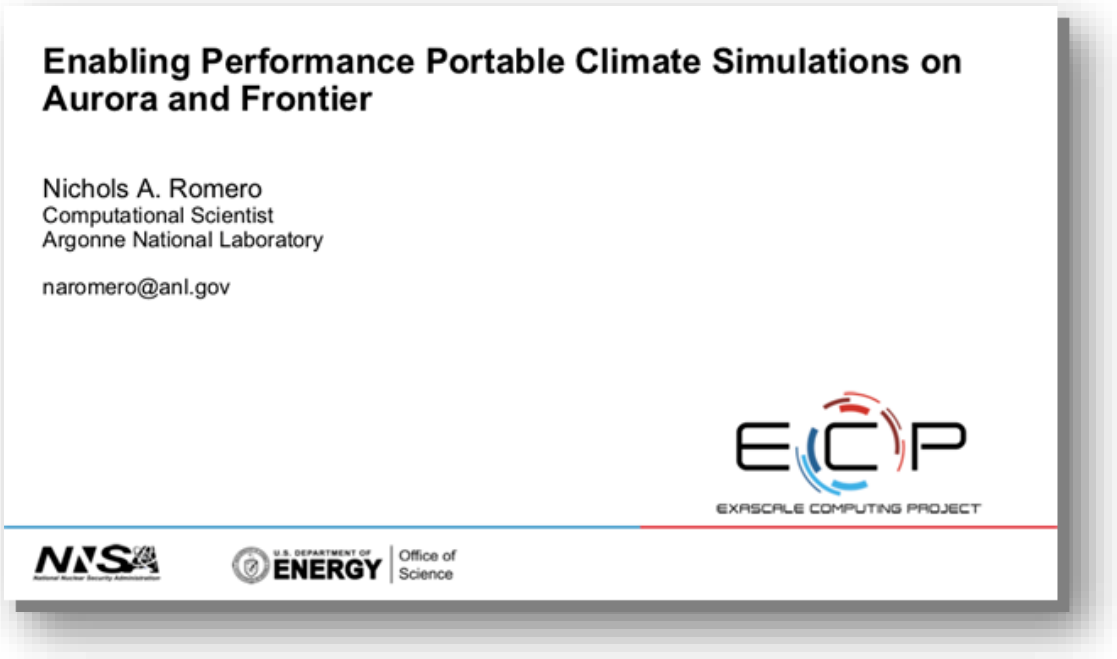

Source: ANL and Hyperion Research, 2020

PRESENTATION: ENABLING PERFORMANCE PORTABLE CLIMATE SIMULATIONS ON AURORA AND FRONTIER, NICHOLS A. ROMERO FROM ARGONNE NATIONAL LABORATORY

Nichols A. Romero described efforts to ensure that global earth system models will run well on DOE pre-exascale computers. These advanced E3SM models include atmosphere, land, ocean, and ice components and are designed to improve DOE's ability to assess regional impacts of climate change on the water cycle that directly affect multiple sectors of the US and global economies, especially agriculture and energy production.

Romero's talk included an overview of E3SM, GPU programming challenges, performance portability methodologies, and project outcomes to date.

FIGURE 1

Source: ANL and Hyperion Research, 2020

Romero described DOE's E3SM as a state-of-the-science Earth system model development and simulation project. E3SM is used to investigate energy-relevant science using code optimized for DOE's advanced computers.

"DOE objectives call for E3SM to provide a 40-year outlook on energy and water issues. While today's systems constrain E3SM to a 10-year outlook, code optimization is ongoing for the accelerated architectures of DOE pre-exascale systems."

FIGURE 2

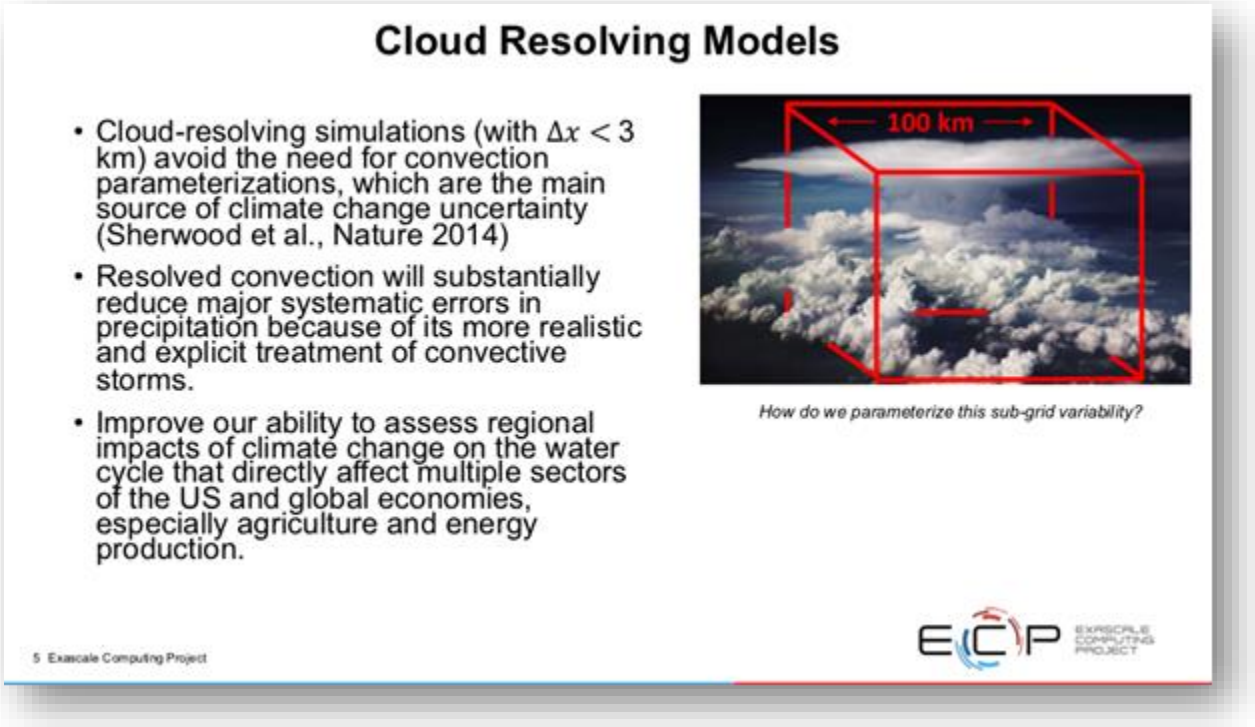

Source: ANL and Hyperion Research, 2020

Romero said that one of the key challenges for E3SM team is to improve its cloud-resolving simulations.

"In our grid, we're trying to simulate the whole earth--the processes going on in the atmosphere, water and land. One of the largest sources of uncertainty in these climate simulations is how we handle the clouds. We want to include convection in clouds, so the more accurately we can we can capture this, the better fidelity our simulations we have. This is a major challenge for us."

Romero described two methods used for cloud-resolving simulations: the Simple Cloud Resolving E3SM Atmosphere Model (SCREAM) and the E3SM-MMF multiscale modeling approach.

- The Simple Cloud Resolving E3SM Atmosphere Model (SCREAM) is a rewrite of existing atmosphere code in C++/Kokkos for GPU support with simplified physics. It scales up to a 3km resolution and has received a 1-year INCITE 1 allocation for producing 10 simulated years of results in 2021.
- The E3SM-MMF Multiscale modeling approach uses "super-parameterization" for cloud resolving convection. This method has high computational density and is ideal for GPUs. Resolved convection will substantially reduce major systematic errors in precipitation because of its more realistic and explicit treatment of convective storms.

For exascale, the E3SM challenge is to further develop this Earth system model with a fully weather-resolving atmosphere and cloud-resolving super-parameterization, an eddy resolving ocean, and ice components, all while obtaining the necessary throughput to run 10-100 member ensembles of 100-year simulations.

FIGURE 3

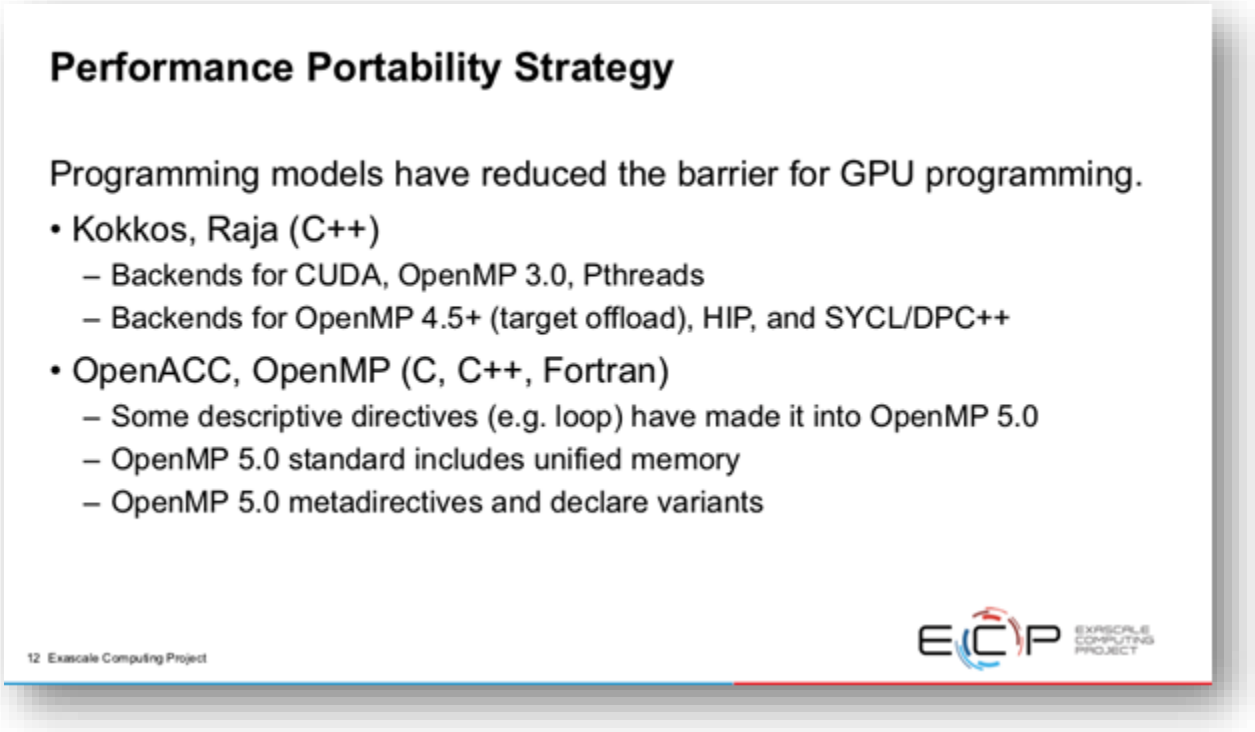

Source: ANL and Hyperion Research, 2020

Romero stated that E3SM codes need to be optimized for GPUs in order to take advantage of the accelerated architectures of DOE's Summit, Frontier, and Perlmutter supercomputers. He went on to describe the E3SM portability strategy.

"Our performance portability strategy is a lot easier now than it was many years ago. We have Kokkos and Raja. These programming models are limited as C++, but it saves us from having to think too hard about whether we are going to need to use CUDA, HIP, SYCL, or something else. It hides that extra level of pain for us."

According to Romero, the Kokkos porting team takes the "brunt" of transition work for new architectures, and E3SM developers will still need to a lot of testing.

Romero described how directives in OpenACC have influenced the development of OpenMP 5.0. The E3SM project will migrate from OpenACC to OpenMP 5.0 by 2021. To accomplish this, the team will

leverage OpenACC-2-OpenMP Source-to-Source translator code, available at: https://github.com/naromero77/ACC2OMP.

FIGURE 4

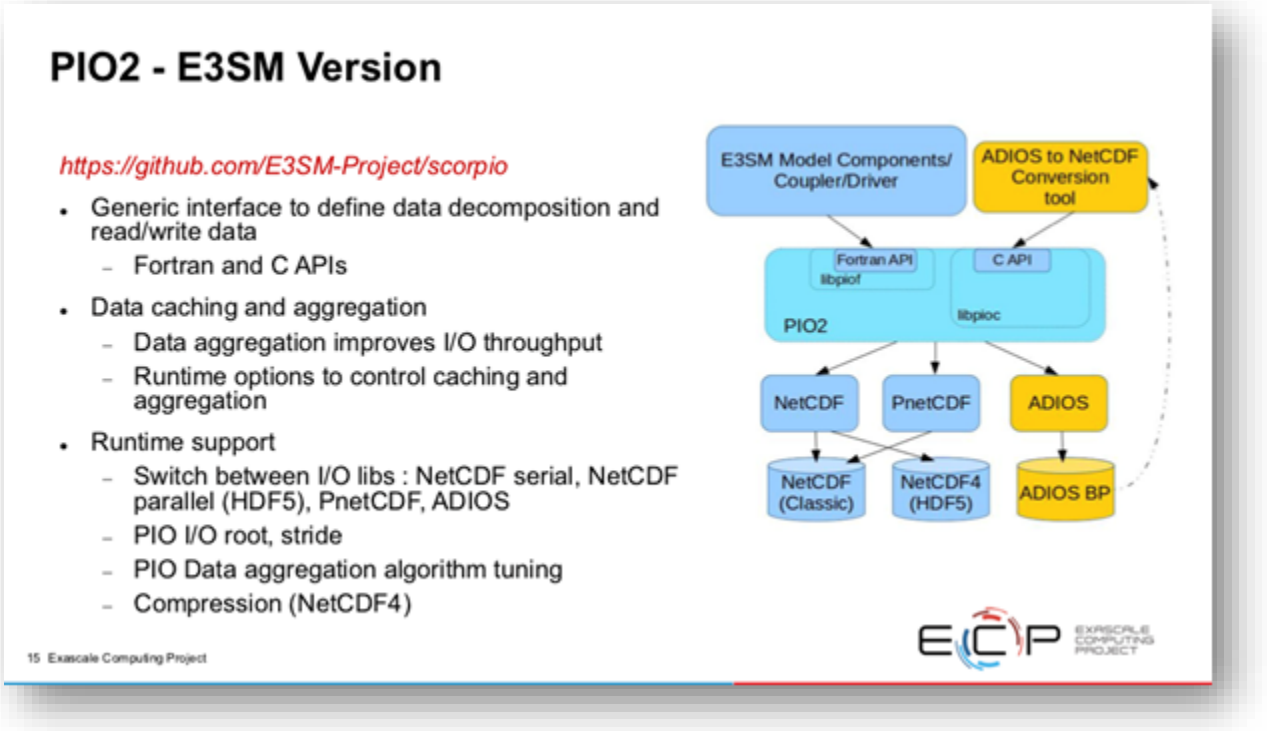

Source: ANL and Hyperion Research, 2020

Romero briefly described encouraging performance results from using the new PIO2 (Scorpio) IO library in their porting efforts. The Parallel Input/Output library (PIO) is used by all the model components in E3SM for reading input and writing model output.

In recent months, developers added support in PIO to read and write data using the Adaptable I/O System (ADIOS) library. The ADIOS library provides a flexible way to describe scientific data that may be read, written, or processed outside a simulation.

Romero summarized key Scorpio performance results:

- With a better caching and aggregation strategy than PIO1, Scorpio delivers up to 10X better IO performance on 96,000 MPI processes on Summit.
- By using ADIOS to write multiple files instead of a single file, Scorpio delivers up to 4X performance compared to PnetCDF implementations.
- Since ADIOS writes data out in multiple files and does not require data to be written out in contiguous chunks, it saves time by partially rearranging data and reducing contention in the file system.

- With more data (higher frequency, longer runs) Scorpio delivers higher throughput for ADIOS, ~140 GB/s.

Conclusions

The E3SM project is on track to achieve its goals, making its earth model ready for pre-exascale systems and beyond. Romero summarized E3SM milestones to date on the Summit supercomputer:

- Completed full port of the Climate Research Model:
  - 20,000 lines of code, using Fortran/OpenACC
  - 98% of the atmosphere cost (reduced radiation configuration)
  - Extensively documented in CAAR final report and ECP Milestone report
- Early Science Award on Summit:
  - Used 350K node hours (1024 nodes)
  - 1km 2D CRM. Completed 7-year simulation
  - Performance and simulations result in Early Science Final Report
- Gordon Bell SC2019 Submission:
  - 1.7km 3D CRM running on 4600 Summit nodes and obtaining 1.8 SYPD
  - 2.5% of peak double precision flops

Romero concluded by inviting attendees to the upcoming DOE Performance, Portability, and Productivity Meeting, April 7 - 9, 2020 in Kansas City, Missouri. Details are available at: https://p3hpcforum2020.alcf.anl.gov/.

For more information or to view this and other presentations given at HPC User Forums dating back to 2008, visit www.hpcuserforum.com.

About Hyperion Research, LLC

Hyperion Research provides data driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). We provide thought leadership and practical guidance for users, vendors and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA
612.812.5798
www.HyperionResearch.com and www.hpcuserforum.com

Copyright Notice

Copyright 2020 Hyperion Research LLC. Reproduction is forbidden unless authorized. All rights reserved. Visit www.HyperionResearch.com of www.hpcuserforum.com to learn more. Please contact 612.812.5798 and/or email info@hyperionres.com for information on reprints, additional copies, web rights, or quoting permission.