

PERFORMANCE ANALYTICS FOR COMPUTATIONAL EXPERIMENTS (PACE)

Understanding computational performance of a complex coupled climate model like the Energy Exascale Earth System Model (E3SM, e3sm.org) poses a singular challenge to domain and computational scientists. Toward that goal, researchers developed PACE (Performance Analytics for Computational Experiments), a framework to summarize performance data collected from E3SM experiments to derive insights and present them through a web portal. The primary goal of PACE is to serve as a central hub of performance data to provide a summary of E3SM experiment performance.

PACE is designed to enable the following capabilities:

- Interactively analysing experiments and application sub-regions,
- Tracking performance benchmarks and simulation campaigns of interest,
- Facilitating performance research on load balancing and process layouts,
- Identifying bottlenecks to inform targeted optimization efforts

WEB PORTAL FEATURES

The PACE portal provides a powerful search capability, including autocomplete, to select and drill

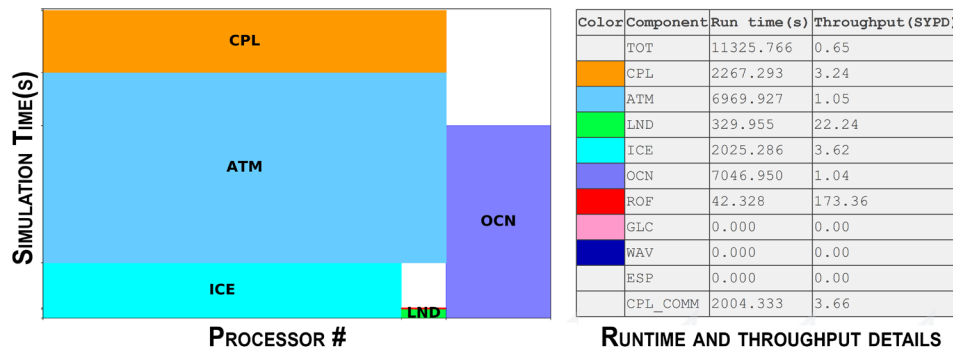


Figure 1. Left: example of a PACE “experiment details view” that shows the mapping of parallel processes to model components along with simulation time spent. Right: model runtime and throughput details listed by component.

down into experiments of interest. Users can search for experiments pertaining to a specific user or platform, or search for a specific model configuration (compset, grid, or case) or a combination thereof. The “search results overview” shows basic experiment metadata such as user, machine, case name, etc. Upon selection of a desired experiment, the user is directed to the “experiment details view” that includes the model throughput and process layout for the various components (see Fig. 1 above).

A user can delve deeply into a particular task’s performance data through an interactive tree graph (see Fig. 2, next page), which displays “application timers” that log how long certain code blocks took to run. Additionally, a user can select a parallel process or thread for a more detailed view aiding in comparisons with other regions or sub-regions.

The flame graph (Fig. 3, next page) illustrates an alternate view of the data that highlights the time spent in different application regions.

PACE also enables performance comparisons across multiple experiments and parallel processes. An interested user can download an experiment’s provenance and raw performance data for further analysis. A detailed demonstration of the portal features is captured in the

Get PACE

- PACE Portal – <https://pace.ornl.gov/>
- Reference Page – <https://e3sm.org/resources/tools/diagnostic-tools/pace>
- Videos –
- Web Portal Features: <https://youtu.be/0U-uu-5hn2Y> and
- How to Upload Data: <https://youtu.be/Vd3q3mJPQfs>

PACE Web Portal Features video
 (<https://youtu.be/0U-uu-5hn2Y>).

METHODOLOGY AND INFRASTRUCTURE

The E3SM model incorporates a lightweight performance tracking capability by default. This tracking/profiling capability is provided by marking various application sections of interest using start and stop markers based on the General Purpose Timing Library (GPTL). Whenever an E3SM user runs an experiment on a supported DOE supercomputer, provenance and performance data is automatically collected and archived in a project-wide location. Such aggregated performance data is periodically uploaded to a central server and made accessible to E3SM users and performance specialists through the PACE web portal.

EXPERIMENT UPLOAD

All members of the E3SM project can upload their experiments' performance data through a streamlined process. The PACE server ingests the performance data from an uploaded experiment to store associated provenance and performance in a database. The raw experiment data is parsed to populate various database tables

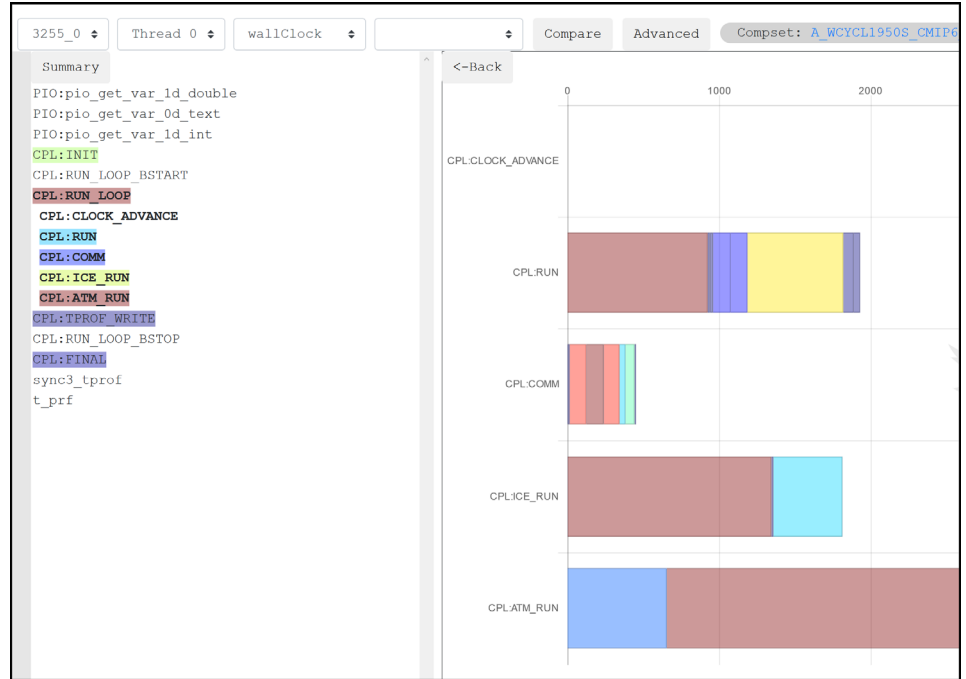


Figure 2. A tree graph view of application performance. Users can click on a specific application timer to dive deeper into the hierarchy to see the timing breakdown in sub-regions of interest.

to facilitate visual analytics and interactive performance exploration.

For instructions on how to upload experiments, see the Upload How-To page (<https://pace.ornl.gov/upload-howto>) and the PACE Upload demonstration video (<https://youtu.be/Vd3q3mJPQFs>).

SUPPORT

This research was supported as part of the E3SM project, funded by the U.S. Department of Energy (DOE),

Office of Science, Office of Biological and Environmental Research.

The students' work was partially supported by an appointment to the Science Education and Workforce Development Programs at Oak Ridge National Laboratory, administered by ORISE through DOE's Oak Ridge Institute for Science and Education.

This research used resources of the Compute and Data Environment for Science (CADES) at the Oak Ridge National Laboratory, which is supported by DOE's Office of Science under Contract No. DE-AC05-00OR22725.

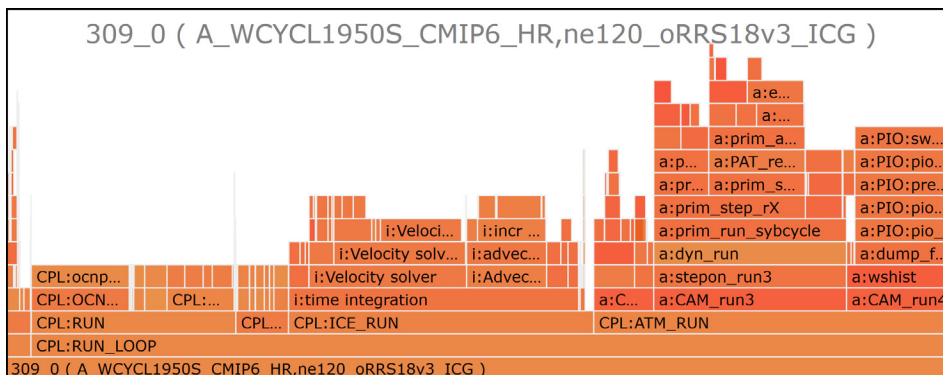


Figure 3. A flame graph summarizes the time spent in various application regions.

CONTACTS

Sarat Sreepathi, Ph.D.
 Oak Ridge National Laboratory
sarat@ornl.gov

Sally McFarlane, Ph.D.
 DOE Program Manager
 Office of Biological and Environmental Research
Sally.McFarlane@science.doe.gov