Managing Computational Experiments

Anshu Dubey (she/her)
Argonne National Laboratory

Software Productivity and Sustainability track @ Argonne Training Program on Extreme-Scale Computing summer school

Contributors: Jared O’Neal (ANL), Anshu Dubey (ANL)

License, Citation and Acknowledgements

License and Citation

- **The requested citation the overall tutorial is:** Anshu Dubey, David E. Bernholdt, Greg Becker, and Jared O'Neal, Software Productivity and Sustainability track, in Argonne Training Program on Extreme-Scale Computing, St. Charles, Illinois, 2023. DOI: 10.6084/m9.figshare.23823822.
- Individual modules may be cited as *Speaker, Module Title*, in *Tutorial Title*, …

Acknowledgements

- This work was supported by the U.S. Department of Energy Office of Science, Office of Advanced Scientific Computing Research (ASCR), and by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration.
- This work was performed in part at the Argonne National Laboratory, which is managed by UChicago Argonne, LLC for the U.S. Department of Energy under Contract No. DE-AC02-06CH11357.
- This work was performed in part at the Lawrence Livermore National Laboratory, which is managed by Lawrence Livermore National Security, LLC for the U.S. Department of Energy under Contract No. DE-AC52-07NA27344.
- This work was performed in part at the Los Alamos National Laboratory, which is managed by Triad National Security, LLC for the U.S. Department of Energy under Contract No.89233218CNA000001
- This work was performed in part at the Oak Ridge National Laboratory, which is managed by UT-Battelle, LLC for the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.
- This work was performed in part at Sandia National Laboratories. Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Running simulations for science discovery is more of a craft and less of science. More than any other aspect of computational science it relies on experience and acquired wisdom that helps one develop a nose for fruitful possibilities.

Why do you need to plan?

- Machines are expensive and rare resources
  - Operating them is also very expensive
  - Many people are competing for these resources
  - You are likely charting new scientific territory

Why do you need to plan?

- Machines are expensive and rare resources
  - Operating them is also very expensive
    - If you made a mistake in your input parameters and got garbage results on a large scale run, you just wasted hundreds of thousands of dollars
  - Many people are competing for these resources
    - Your wasted run is likely to be either your or someone else's opportunity lost
  - You are likely charting new scientific territory
    - Some aspect of using your code may not have been important before, but may become critical in the new study
    - Some solver may run up against the limits of its validity
    - Inflight correction may be needed to parameters to continue with the study
- Aim for no surprises, but be prepared for them

Why do you need to plan?

- Machines are expensive and rare resources
  - Operating them is also very expensive
    - If you made a mistake in your input parameters and got garbage results on a large scale run, you just wasted hundreds of thousands of dollars
  - Many people are competing for these resources
    - Your wasted run is likely to be either your or someone else's opportunity lost
  - You are likely charting new scientific territory
    - Some aspect of using your code may not have been important before, but may become critical in the new study
    - Some solver may run up against the limits of its validity
    - Inflight correction may be needed to parameters to continue with the study
- Aim for no surprises, but be prepared for them

In the 2005 simulation mentioned earlier, out of 5 teams, ours was the only team that had success in getting a good science outcome

How do you plan

- Focused verification of the target simulation on the target platform
  - Over and above regular testing
  - Emphasis on understanding solver validity regime

- Pathfinder runs to get a good estimate of needed resources
  - Cost benefit analysis of fidelity vs reaching science goals in allocated resources

How do you plan

- Develop helpful diagnostics
  - Low overhead ways of confirming the health of the run
    - Are conserved quantities conserved?
    - Has any quantity become unphysical?

- Develop hierarchy of analysis
  - Full analysis of runs is not feasible in flight
  - Intermediate level analysis can give further insight into health of the simulation

Story of one simulation campaign

- Theory of Type Ia supernova explosion – 2006/2007
  - Evidence from observations:
    - Light curve powered by Ni56 decay
    - Evidence of medium weight elements, but in much smaller quantities
  - Implied transition from deflagration to detonation
- A 2D exploratory run had given a tantalizing answer to how?
  - To confirm a full 3D run was needed at good enough resolution
  - It would be the largest run of its kind at the time – totally uncharted territory
  - Until then 3D runs had been octants relying on symmetry
    - The 2D run had shown that symmetry had to be avoided

Preparation Steps

- Step 1 – develop a test that represents the most complex physics interactions
- Challenges:
  - Features take a long time to develop
  - Want to ensure that at least one refinement step occurs during the test
  - IO too slow to restart from a large checkpoint at late stage of the run
    - Also test would need a large chunk of the machine
- Use physics understanding to create initial conditions that would quickly develop comparable complexity

Preparation Steps

- Step 2 – Use the new test to characterize the performance behavior of the target platform
- Motivation:
  - Standard performance studies could not give crucial information
  - AMR refinement patterns make each application different
  - Interoperability and trade-off opportunities needed to explored in a closely resembling simulation behavior
- Full fidelity 2D runs, and a set of runs of the new test provided enough information to extrapolate and estimate needed CPU hours

Preparation Steps

- Step 3 – Look for trade-offs and optimization opportunities
- Motivation:
  - Initial CPU estimates too high to complete the runs within allocations
  - Exploration of any parameter space needs to minimize individual run times
- Many opportunities were found, documented in reference below.
  - Identify redundant refinement and get rid of it
  - Coarsen computations for some physics
  - Move some computations to post-processing
- **All optimizations were based on scientific and numerical intuitions**

Dubey A, Calder AC, Daley C, et al. Pragmatic optimizations for better scientific utilization of large supercomputers. The International Journal of High Performance Computing Applications. 2013;27(3):360-373. doi:10.1177/1094342012464404

Preparation Steps

- Step 4 – Prepare diagnostics and quick analysis mechanism
- Examples-- diagnostics
  - Conservation of mass, momentum energy
  - Changes in dt recorded in the logfiles
  - Spikes in variable values
- Examples – quick analysis
  - Quick visualization of random 2D slices
  - Inspection of critical quantities in 1D

Preparation Steps

- Step 4 – Prepare diagnostics and quick analysis mechanism
- Examples-- diagnostics
  - Conservation of mass, momentum energy
  - Changes in dt recorded in the logfiles
  - Spikes in variable values
- Examples – quick analysis
  - Quick visualization of random 2D slices
  - Inspection of critical quantities in 1D

Lab notebook artifacts

- every run registers all configurations and runtime parameters in a logfile.
- logfiles are cumulative
- dedicate space for storing all results in a preconfigured directory structure
- scripts to move output from scratch to the dedicated space

Outcome

- A successful campaign
  - But not without hitches
- Optimization related runs were not given the same level of care
  - Data was considered disposable
  - Code changes were documented

Outcome

- A successful campaign
  - But not without hitches
- Optimization related runs were not given the same level of care
  - Data was considered disposable
  - Code changes were documented
- For the paper the referee asked for details from optimizations
  - We did not have them
  - Fortunately the referee was satisfied with reasoning and other supporting evidence we produced

Summary and Takeaways

- Good science with computation is a craft -- training is needed in how to do it
- Machines are expensive to build and expensive to run
  - They provide opportunity for great work
  - Care is needed to ensure that the outcome meets expectations
- Reproducible results are a necessity, not a luxury
  - There is no credible science without provenance

"a parameter combination that induces erroneous results is easily selected“
\- https://doi.org/10.1063/1.476021