PERFORMANCE, POWER, AND ENERGY OF IN-SITU AND POST-PROCESSING VISUALIZATION: A CASE STUDY IN CLIMATE SIMULATION

Vignesh Adhinarayanan*, Scott Pakin†, David Rogers†, Wu-chun Feng*, James Ahrens†

* Department of Computer Science, Virginia Tech, Blacksburg, VA24060

† CCS-7 Division, Los Alamos National Laboratory, Los Alamos, NM 87544

Introduction

- Off-chip data movement can consume hundreds of times as much energy as on-chip data movement
- More data produced from high-resolution simulation to increase fidelity ➔ More power/energy for storage subsystem
- Problematic because future supercomputers will be power limited

Operation	Energy (pJ)
DF FLOP	10
Register	1
1mm on-chip	3-5
5mm on-chip	20
Off-chip	1000-2000

Energy consumption projection for an exascale system [1]

Hypothesis

Reducing disk reads and writes using the following techniques will save significant amount of energy and power:

- Temporal sampling – Write output only every few time steps
- In-situ visualization – Produce images *during* simulation (without writing raw data to the disk) and write only the compact image representation

Experimental Setup

Hardware Details

Component	Detail
Processor	2x Intel Xeon E5-2665 @ 2.4GHz
DRAM	4x 16GB DR3-1333
Disk	500GB Seagate 7200rpm

Power Measurement

Power measured at 1 Hz frequency using the following methods for different components:

- **Full system** – WattsUp Pro power meter
- **Processor and DRAM** – Intel RAPL interface (statistical model based on performance counters)
- **Disk** – Statistical power model based on *iostat* statistics

Application

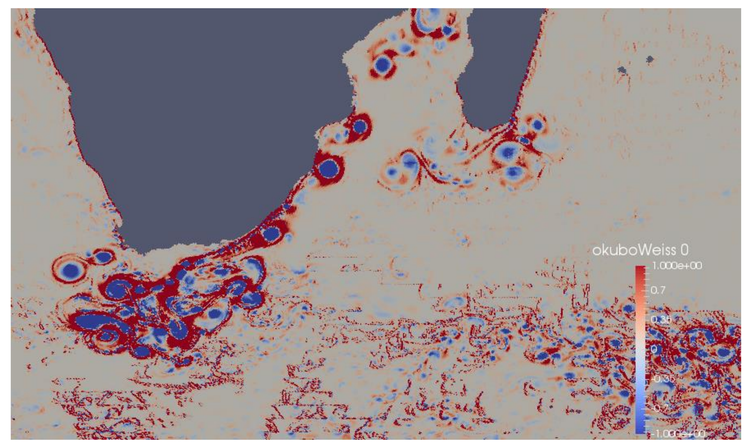

MPAS Ocean simulation

Ocean component of the modeling for prediction across scale (MPAS-O) [2] solves an unstructured mesh problem to calculate the Okuba-Weiss metric. The end goal is to identify eddies in the ocean (shown in figure). Visualization through Paraview-Cinema [4].

Problem Size: 240km grid run for simulated period of one month

Disk power modeling

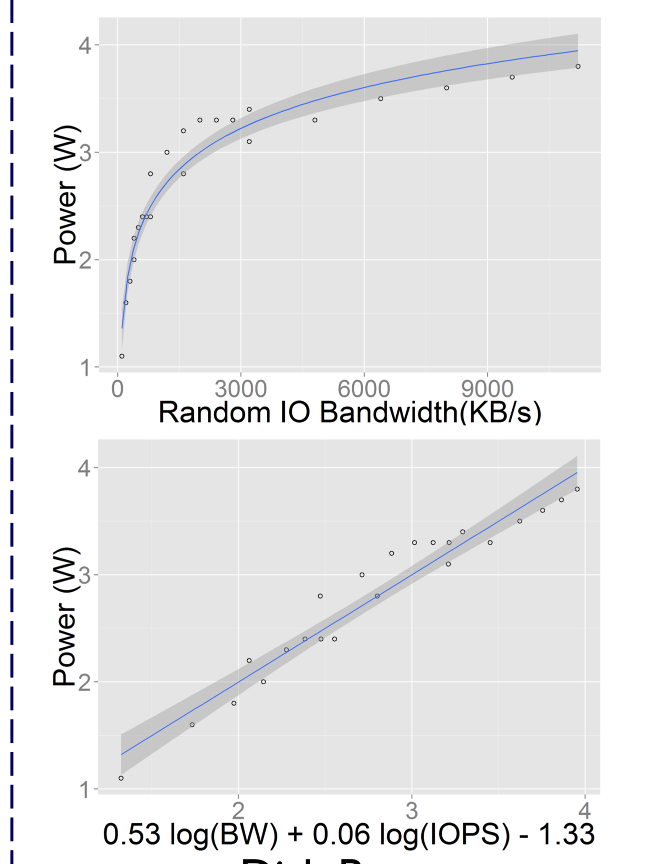

Disk Power

5.67 + 0.53*log(BW) + 0.06*log(IOPS)

Results

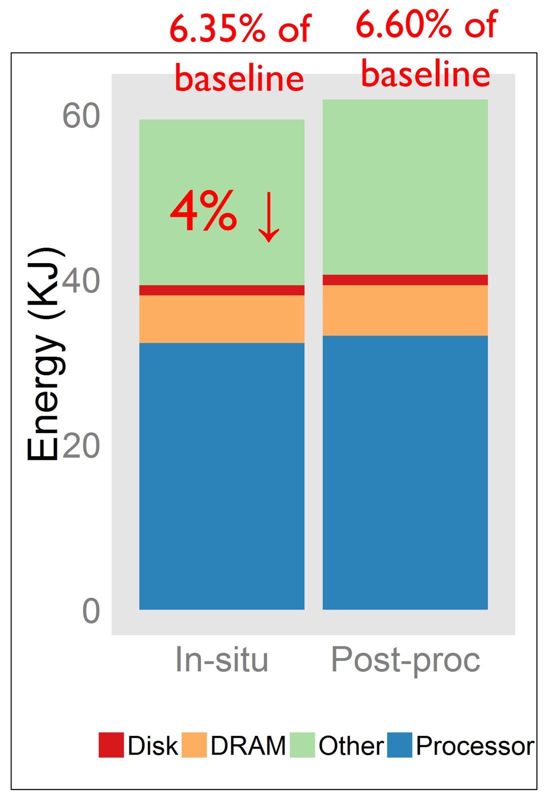

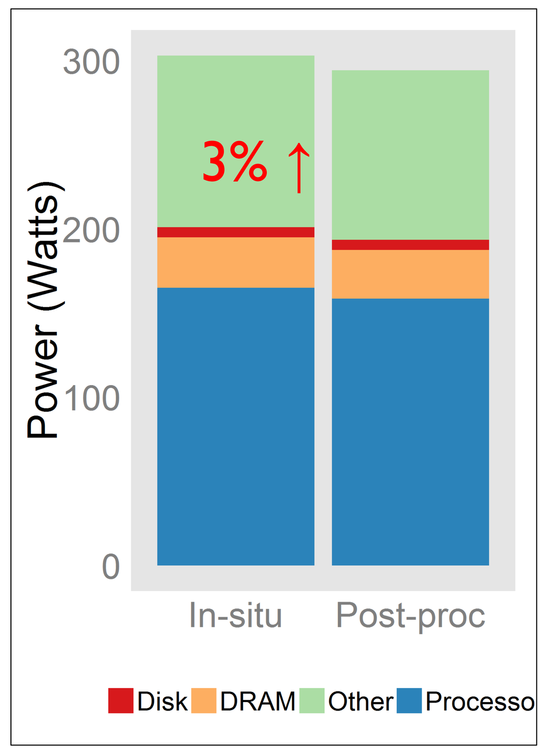

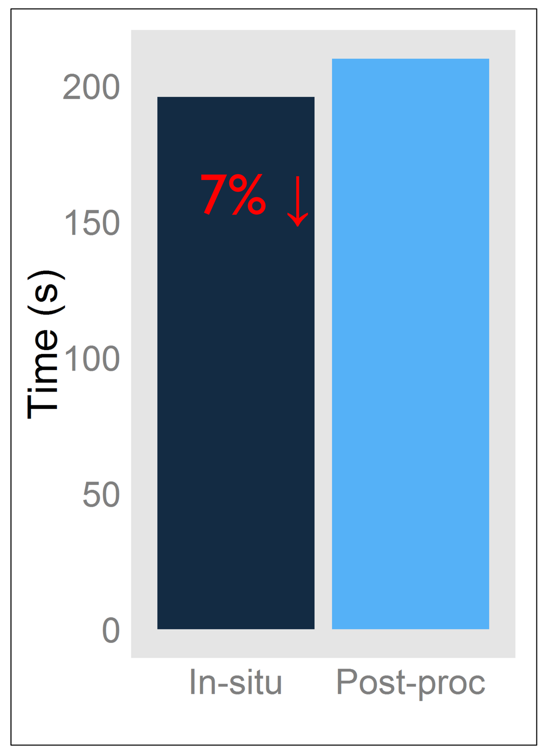

- Baseline – "Traditional" post-processing *without* any sampling
- Post-processing – "State-of-the-practice" post-processing with *temporal sampling* (i.e., write every few iterations – in this case 24 iterations)
- In-situ – Produce images *in situ* alongside simulation and write compact image representation once every 24 iterations)
- In-situ visualization saves 4% energy for MPAS-O for the given problem size
  - Despite consuming 3% more power on an average
  - Amortized by 6.7% lower execution time from reduced I/O wait time
- Energy saved from disk subsystem almost negligible
  - Nearly all energy saved from reduced system idling
- 97.5% lower storage requirement for in-situ pipeline
- Preliminary results on a 128-node cluster: 55% energy savings for in-situ pipeline (for 60km grid size, 1output/simulated day sampling rate)

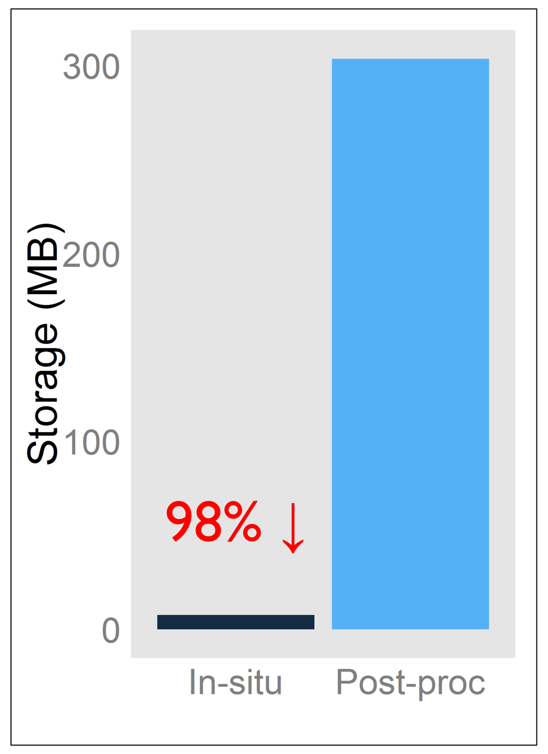

Discussion

- Lower storage requirements ➔ Fewer I/O nodes
- Fewer I/O nodes ➔ More power budget for compute nodes
  - Assuming 10% nodes reserved in a HPC data center for storage and assuming storage nodes consume same power as compute nodes, data center power goes down by nearly 10%
  - Estimated increase in power budget for compute nodes ~10%
- 10% power budget increase shows up to 6.3% improvement in performance for MPAS-O using RAPL interface
- Future work: Burst buffers offer opportunity to reduce I/O wait time, but at the cost of increased power consumption by the burst buffers

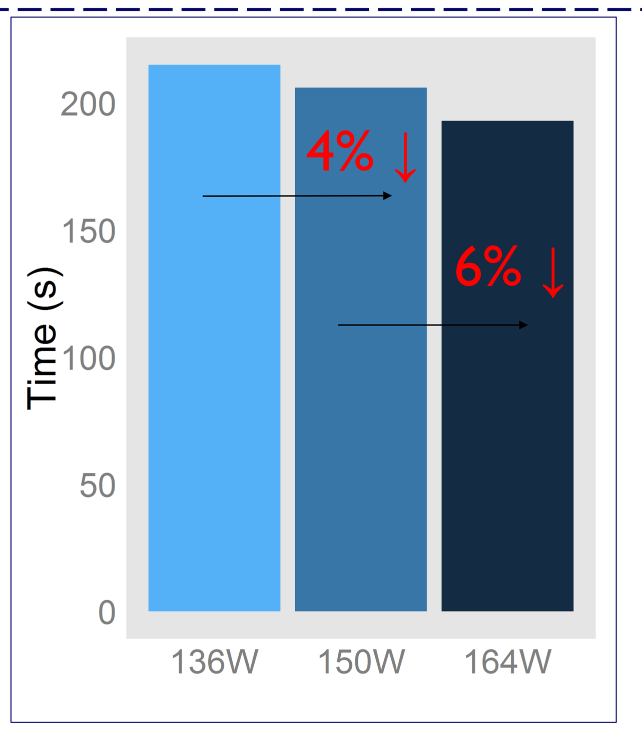

Conclusion

In-situ visualization offers the following advantages:

- Reduced energy consumption (by reducing system idling or I/O wait time)
- Reduced power (by using fewer storage nodes)
- Improved performance (by reducing I/O wait time and by making more power available for compute nodes)

Bibliography

[1] S. R. Sachs, K. Yelick et al., "Exascale Programming Challenges," *2011 Workshop on Exascale Programming Challenges*, 2011.

[2] T. Ringler et al., "A Multi-Resolution Approach to Global Ocean Modelling," *Ocean Modelling, 69(C), 211–232.*

[3] V. Adhinarayanan et al., "On the Greenness of In-situ and Post-Processing Visualization Pipelines," *Proceedings of the 11th Workshop on High-Performance Power Aware Computing (HPPAC)*, May 2015.

[4] Ahrens et al., "An Image-based Approach to Extreme Scale in Situ Visualization and Analysis," Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)", Nov 2014

Acknowledgment

This work was supported in part by a grant from the U.S. Department of Energy (DOE) Office of Advanced Scientific Computing Research (ASCR) via DE-SC0012637 and equipment donation from Supermicro. The authors wish to thank Greg Abram for his feedback on this work.

Virginia Tech Invent the Future — SyNeRG — Los Alamos National Laboratory EST. 1943