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)
DP 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

Single-Node Setup

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

Disk power modeling

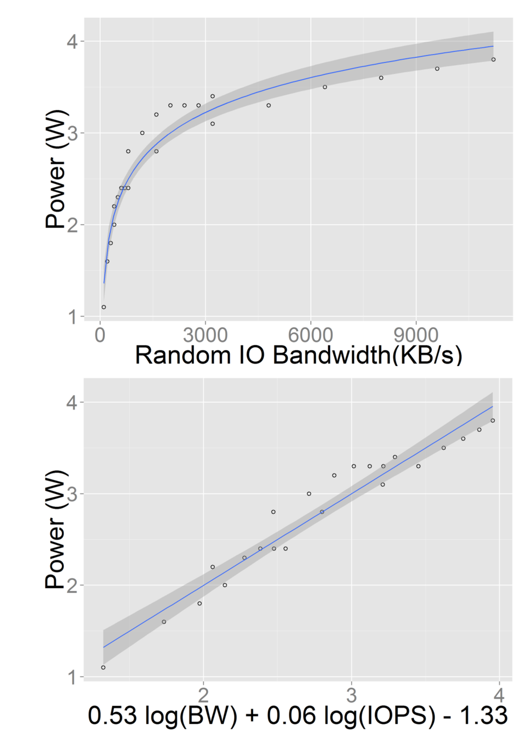

Disk Power
5.67 + 0.53*log(BW) + 0.06*log(IOPS)

HPC System Setup

- Compute cluster
  - 128 nodes of Caddy supercomputer
  - 2x Intel E5-2670 CPU/node
  - 64 GB RAM/node
  - Power measured for 10 nodes using cage power meter and extrapolated
- Storage cluster
  - 5 nodes running Lustre file system
  - 1 master node, 2 metadata servers, 2 object storage servers
  - Intelligent PDUs for power measurement

Application

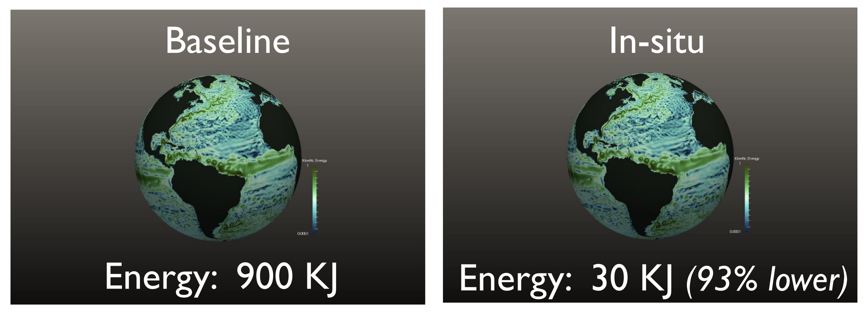

Same cognitive value for both visualization pipelines

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: 240-km grid run for simulated period of one month

Results

Visualization Pipelines Evaluated

1. Baseline – "Traditional" post-processing *without* any sampling
2. Post-processing – "Modern" post-processing with *temporal sampling* (i.e., write every *n* iterations – in this case, *n* = 24)
3. In-situ – Produce images *in situ* alongside simulation and write compact image representation once every 24 iterations)

Key Findings

1. In-situ Visualization vs. Baseline ("Traditional" Post-Process)
   - *Saves 93% energy* for MPAS-O for the given problem size
     - … despite consuming 3% more power on average
     - … but amortized by 94% faster execution from reduced I/O wait
2. In-situ Visualization vs. Post-processing ("Modern" Post-Process)
   - *Saves 4% energy* for MPAS-O for the given problem size
     - … despite consuming 3% more power on average
     - … but amortized by 7% faster execution from reduced I/O wait
3. Energy saved from disk subsystem almost negligible
   - *Nearly all energy saved from reduced system idling*
4. *97.5% lower storage requirement for in-situ pipeline*

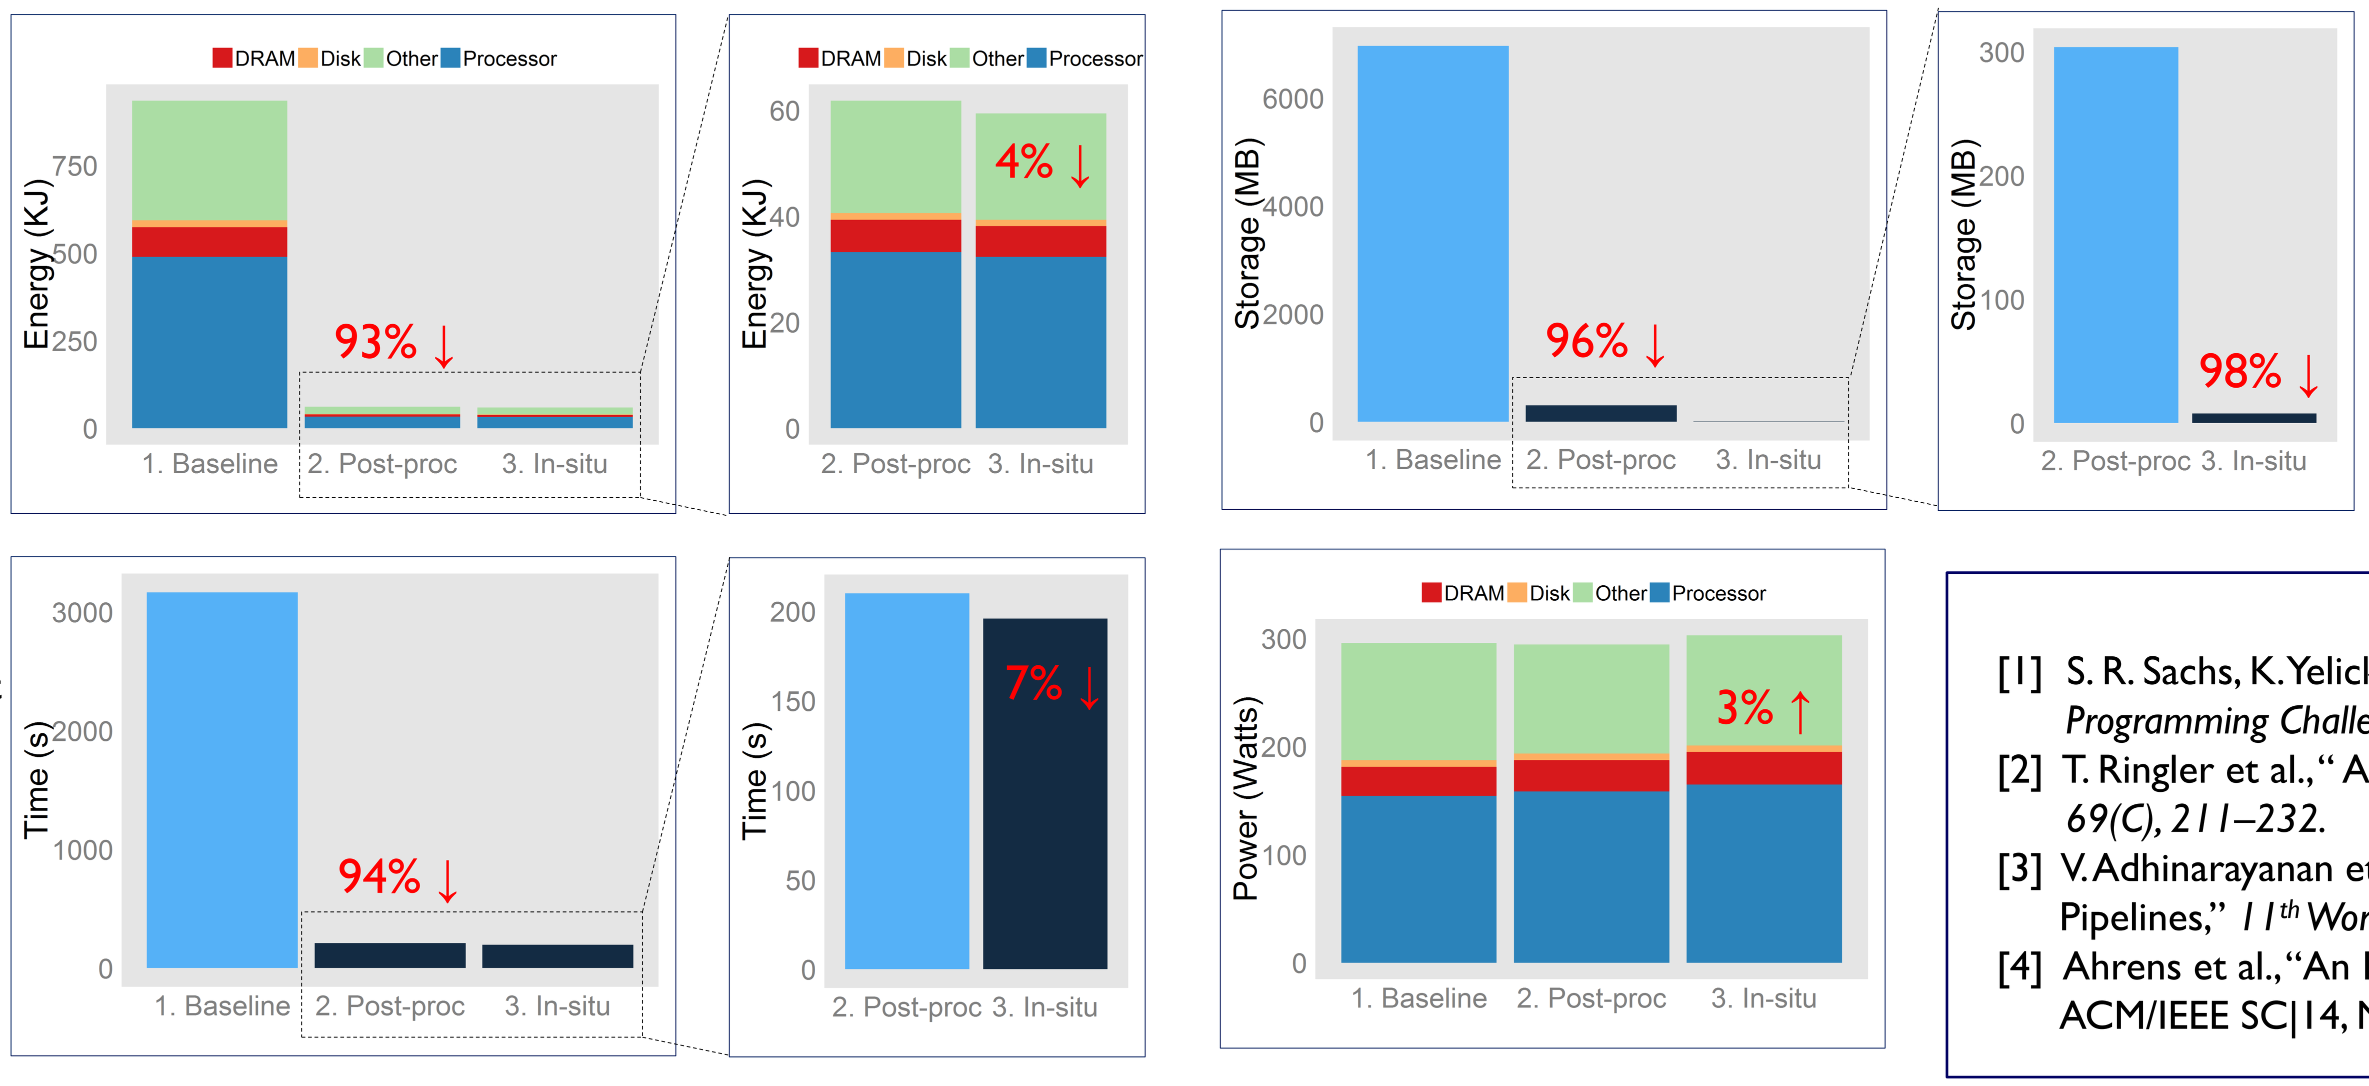

Preliminary Results at Scale

- Problem size: 60-km grid size
- Sampling rate: *One* output per simulated day
- *Key finding: 55% energy savings for in-situ pipeline (vs. modern post-processing pipeline)*
- More aggressive sampling possible to save more energy, but risks missing important events of simulation

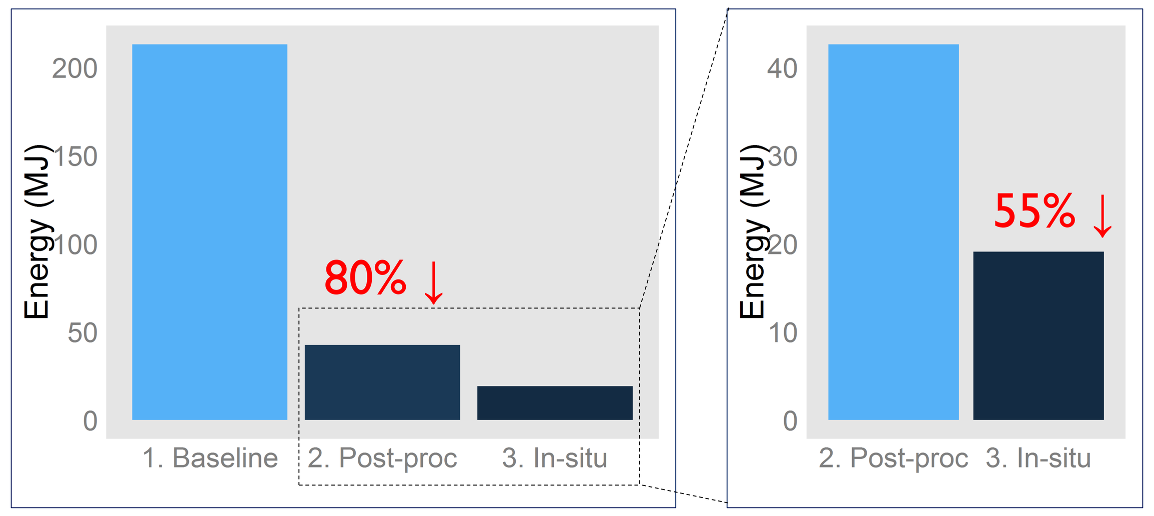

Implications

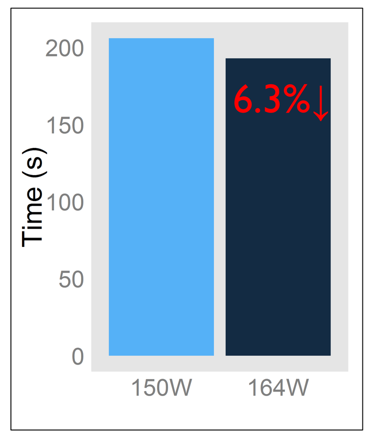

- Lower storage requirements ➔ ***Fewer I/O nodes***
- Fewer I/O nodes ➔ ***More power for compute nodes***
  - Assuming 10% nodes reserved in a HPC data center for storage,
    ➔ ***data center power goes down by ~ 10%***
  - Estimated increase in power budget for compute nodes ~ 10%
    ➔ ***6.3% improvement in performance for MPAS-O using RAPL interface***

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," *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," ACM/IEEE SC|14, Nov 2014.

Acknowledgment

This work was supported in part by Dr. Lucy Nowell, Program Manager for the Advanced Scientific Computing Research (ASCR) program office in the Department of Energy's (DOE) Office of Science via DE-SC0012637. The authors also wish to thank Francesca Samsel for the visualization and Greg Abram for early discussions on this work.

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

This poster is a Los Alamos Unclassified Release LA-UR-15-26284

ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, Austin, TX, Nov. 2015.