Increasing Scientific Data Insights About Exascale Class Simulations Under Power and Storage Constraints

James Ahrens, ahrens@lanl.gov, Los Alamos National Laboratory, A White Paper for the 2nd Big Data Exascale Computing Workshop, February 2014, LA-UR-14-20859

The purpose of this white paper is to frame the challenging problems of exascale data analysis and visualization. A successful framing of these challenges will guide future research and development and highlight the differences from today's approaches. One key difference is the notion of a ***cost per insight*** in terms of power and storage used. This notion of constraints on our insights challenges the premise of our traditional workflow.

Power constraints are driven by reducing the many financial costs, including facility, power and cooling costs, associated with the massive power requirements that are projected for an exascale machine without research and development interventions. Specifically, the United States Department of Energy's exascale strategy identifies target goals for peak performance to increase three orders of magnitude while system power is only targeted to increase by a factor of two. To keep within an extremely limited power budget, locality during computation is extremely important. The most expensive operation is data movement, from both the power and performance perspective, moving data up from the CPU out through the memory hierarchy including out to persistent storage and the network.

Storage constraints are also driven by financial costs including power costs. Future storage technology projections suggest that the gap between both capacity/bandwidth and FLOPS will widen as we move towards exascale. Therefore, we expect the storage system of an exascale supercomputer to be smaller and slower compared in a relative way with the peak FLOPS of today's generation of supercomputers for a proportionally similar level of investment.

In light of these constraints, it is instructive to review current and proposed approaches. In a traditional post-processing oriented visualization and analysis approach, temporal simulation snapshots are saved at regular intervals. This approach incorporates the process of saving checkpoints for later restart in case of errors. Traditionally, these full simulation checkpoint snapshots and additional smaller visualization and analysis data are interactively analyzed after the simulation run is complete. The visualization and analysis community has identified this approach unworkable at extreme scale due to power and storage constraints. An emerging consensus is that significantly more visualization and analysis should occur *in situ*, that is, during the simulation run while the data is resident in memory.

This change of focus from post-processing to *in situ* analysis suggests a set of emerging guidelines about the simulation analysis process:

Sampling and Uncertainty Quantification of Simulation Data are Needed – During *in situ* data analysis the analyst has access to the entire simulation data in all its complexity, including spatial, temporal, multivariate and variable type domains. This data is available only briefly at simulation runtime when it is resident in memory and then deleted when the simulation advances. Given our budgeting constraints, it becomes clear that *in situ* analysis is a form of sampling. The traditional workflow samples fully on the spatial, multivariate and variable type domains at the expense of sampling fully in the temporal domain. Simulation scientists have the opportunity to significantly increase the quality of their analysis results by choosing how to sample from each domain. The quality of their results can be measured through combined *in situ* sampling/uncertainty quantification techniques. For example, in our work, we statistically sample using a stratified random sampling approach on the MC^3 cosmological particle simulation. We store these samples in a level-of-detail organization for later interactive progressive visualization and feature analysis. By sampling during the simulation, we are able to analyze the entire particle population to record full population statistics and quantify sample error [1].

Deliberate Analysis Choices Are Necessary – In the traditional approach, during a simulation run, full simulation snapshots are saved. This has led to the belief that these snapshots can answer arbitrary analysis post-processing questions because "all the data has been saved". As noted above this is not necessarily true for the time domain. A related belief about *in situ* techniques is that automatic selection of data at runtime reduces the type of questions that can be asked about the data during post-processing analysis. It is important to appreciate the traditional post processing approach of saving full simulation snapshots is, in and of itself, an inherently *in situ* activity. Saving full simulation snapshots in time is simply one choice among many for extracting data and/or information from running simulations. An alternative perspective is since our analysis is constrained by a power and storage budget it is important to make deliberate analysis choices before the simulation is run, about what scientific questions will be answered, and then to explicitly save the appropriate data. In the observational/experimental community, pre-planned data reducing streaming analysis is common practice. Custom software and hardware accelerators are typically employed to reduce and analyze data in real-time for accelerator physics, fusion reactors and cyber-security. Our focus on *in situ* approaches aligns the supercomputing community with the observational/experimental community supporting synergistic approaches in the future. Key research questions to answer are: How general and with what quality can analysis questions be answered from compact data products generated *in situ* after the simulation run has completed, in a post-processing manner? What new mathematical or analysis techniques will support this process?

Data Reduction and Prioritization Is Required – An additional corollary to notion of a budget is the requirement that the simulation data stream must be significantly reduced into a compact analysis product in order to fit within the budget. This reduction does not have to be via a statistical sampling; visualization operations and feature extraction algorithms can also be considered a type of sampling strategy. An interesting way to approach the inclusion of the most important data within a budget is to prioritize data using a greedy algorithm saving the highest priority information as the simulation progresses. For example, in recent work we measured temporal entropy in a running simulation. A memory buffer collected time steps with the highest entropy by having time steps with higher entropy overwrite ones with lower entropy. The resulting collection of high entropy time steps provides a summary of the phases of the simulations in which the most change occurs.

[1] Woodring, Jonathan, J. Ahrens, J. Figg, Joanne Wendelberger, Salman Habib, and Katrin Heitmann. "In-situ Sampling of a Large-Scale Particle Simulation for Interactive Visualization and Analysis." In Computer Graphics Forum, vol. 30, no. 3, pp. 1151-1160. Blackwell Publishing Ltd, 2011.