SOSflow: A Scalable Observation System for Introspection and In Situ Analytics

Chad Wood
University of Oregon
Eugene, OR, United States
cdw@cs.uoregon.edu

ABSTRACT

The performance of HPC simulation codes is often tied to their simulated domains; e.g., properties of the input decks, boundaries of the underlying meshes, and parallel decomposition of the simulation space. A variety of research efforts have demonstrated the utility of projecting performance data onto the simulation geometry to enable analysis of these kinds of performance problems. However, current methods to do so are largely ad-hoc and limited in terms of extensibility and scalability. Furthermore, few methods enable this projection online, resulting in large storage and processing requirements for offline analysis. We present a general, extensible, and scalable solution for in-situ (online) visualization of performance data projected onto the underlying geometry of simulation codes. Our solution employs the scalable observation system SOSflow with the in-situ visualization framework ALPINE to automatically extract simulation geometry and stream aggregated performance metrics to respective locations within the geometry at runtime. Our system decouples the resources and mechanisms to collect, aggregate, project, and visualize the resulting data, thus mitigating overhead and enabling online analysis at large scales. Furthermore, our method requires minimal user input and modification of existing code, enabling general and widespread adoption.

About the author:

I am Chad Wood, a fourth-year Computer and Information Science PhD student at the University of Oregon. Prior to my academic pursuits, I enlisting in the U.S. Army in 2005, and was assigned to 2nd Batallion, 75th Ranger Regiment for the duration of my time in service, eventually receiving promotion to the Non-Comissioned Officer rank of Sergeant, and an honorable discharge. I earned a B.A. in Philosophy from Texas Christian University in 2014.

My research focus is on monitoring, introspection, feedback, and control for HPC systems, emphasizing online in situ operations and scalability. I explore this research space using a novel runtime system named SOSflow that I developed at the University of Oregon. SOSflow provides a flexible, scalable, and programmable framework for observation, introspection, feedback, and control of HPC applications.

The Scalable Observation System (SOS) performance model used by SOSflow allows a broad set of online and in situ capabilities including remote method invocation, data analysis, and visualization. SOSflow can couple together multiple sources of data, such as application components and operating environment measures, with multiple software libraries and performance tools, efficiently creating holistic views of performance at runtime.

KEYWORDS

sos, sosflow, hpc, performance, visualization, in situ

REFERENCES