Analyzing Performance at Exascale Using Intuitive Visualizations

Aaditya G. Landge, Katherine E. Isaacs, Abhinav Bhatele, Avishek Saha, Joshua A. Levine, Todd Gamblin, Martin Schulz, Peer-Timo Bremer, Bernd Hamann and Valerio Pascucci

SCI X 2012

Introduction

Supercomputers are increasingly used to perform simulations that model complex scientific phenomenon. Due to the complexity of present simulations and architectures, understanding and analyzing the performance of these simulations is difficult and complex. We present visualizations that help in understanding the performance using intuitive visualizations that are effective in identifying performance bottlenecks and giving better insights.

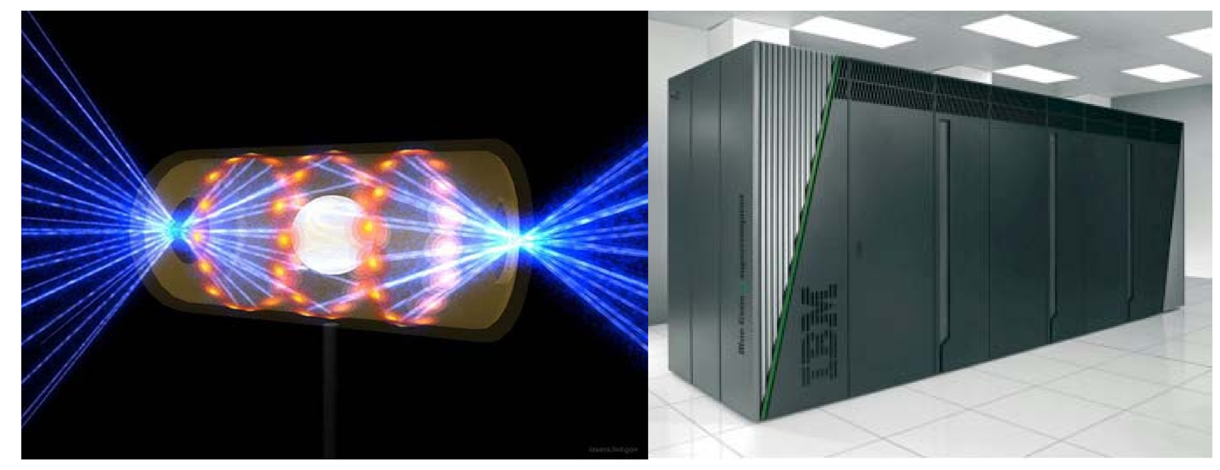

Multi-beam laser plasma simulation (left) IBM BlueGene/Q supercomputer (right)

Philosophy – HAC model

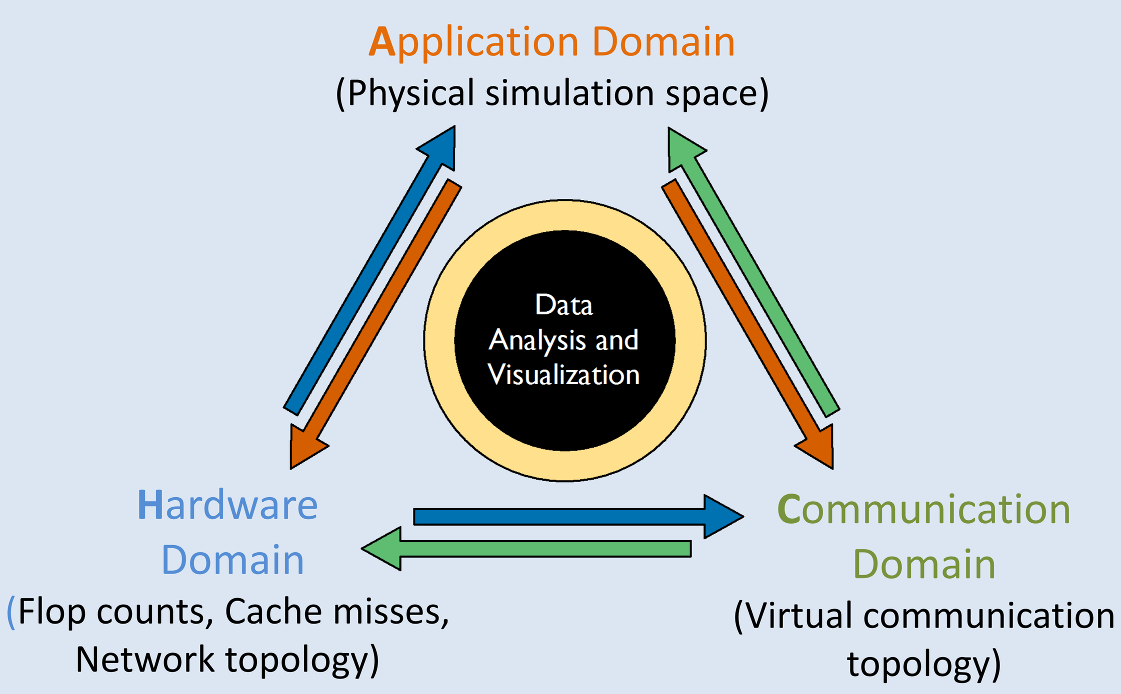

Performance data is collected in three domains: application, hardware and communication. For each domain, we define projections that allow the data to be viewed in the other domains. This framework is called the HAC model. The model not only allows us to directly compare the data across domains, but also allows us to use data visualization and analysis tools available in the other domains.

Projections from Hardware Domain onto the Application Domain

By projecting performance data on the domain of the application one is able to relate the output of the application or the physical simulation space to the performance data giving better insight into the behavior of the system with respect to the application.

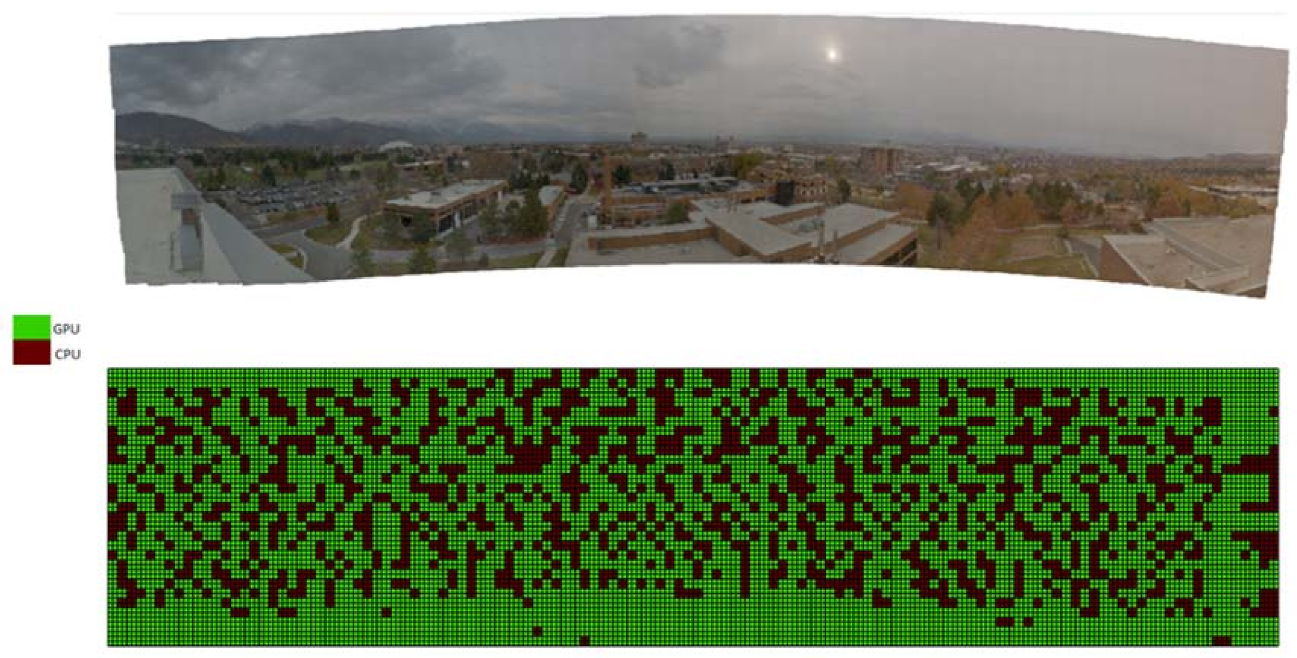

Gigapixel image (top) being processed on a heterogeneous cluster having CPUs and GPUs. Parts of image being processed on CPUs and GPUs (bottom)

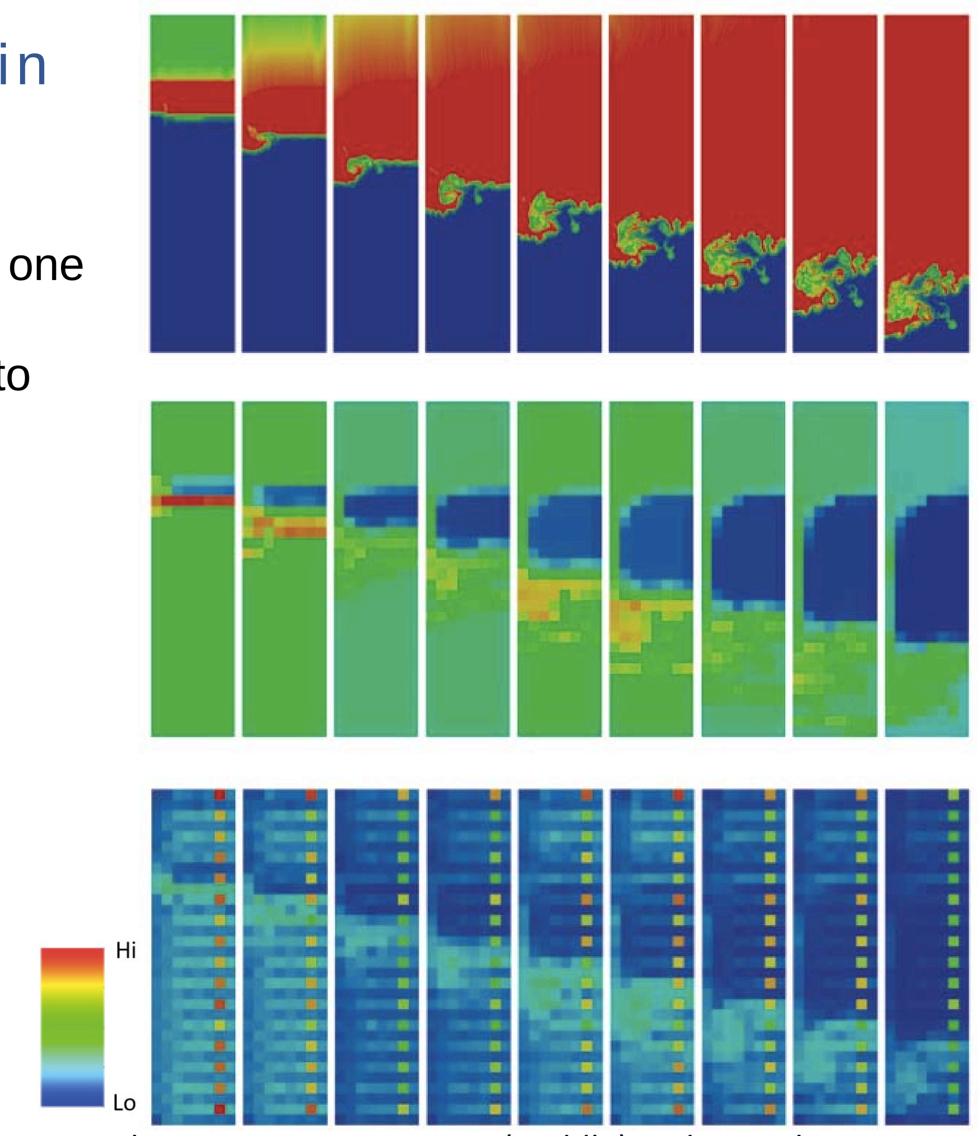

Floating point operations (middle) and L1 cache misses (bottom) mapped to the application domain (top) for multiple time steps of Miranda running on 256 cores of an Infiniband cluster

Projections of Communication domain onto the Hardware Domain

Communication is becoming the dominant performance bottleneck as we scale to a large number of cores. It becomes important to analyze communication in terms of contention on specific links (hot-spots) and distribution of network traffic on the links in various directions.

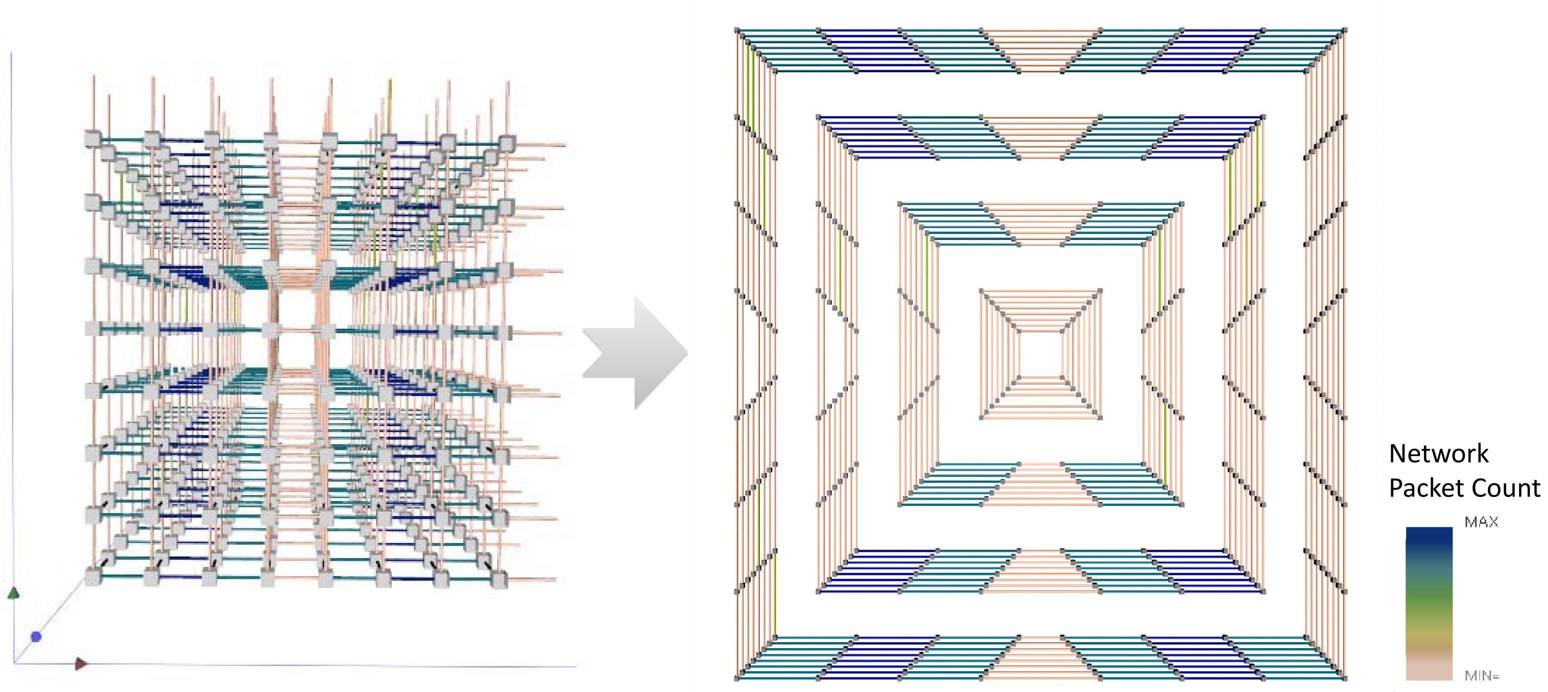

3D torus networks of supercomputers being represented in the 3D view (left). Each compute node represented as a cube connected to other nodes. Edges colored on the basis of network packet counts. Planar projection of the 3D torus (left) shows network traffic without occlusions and reveals the communication patterns.

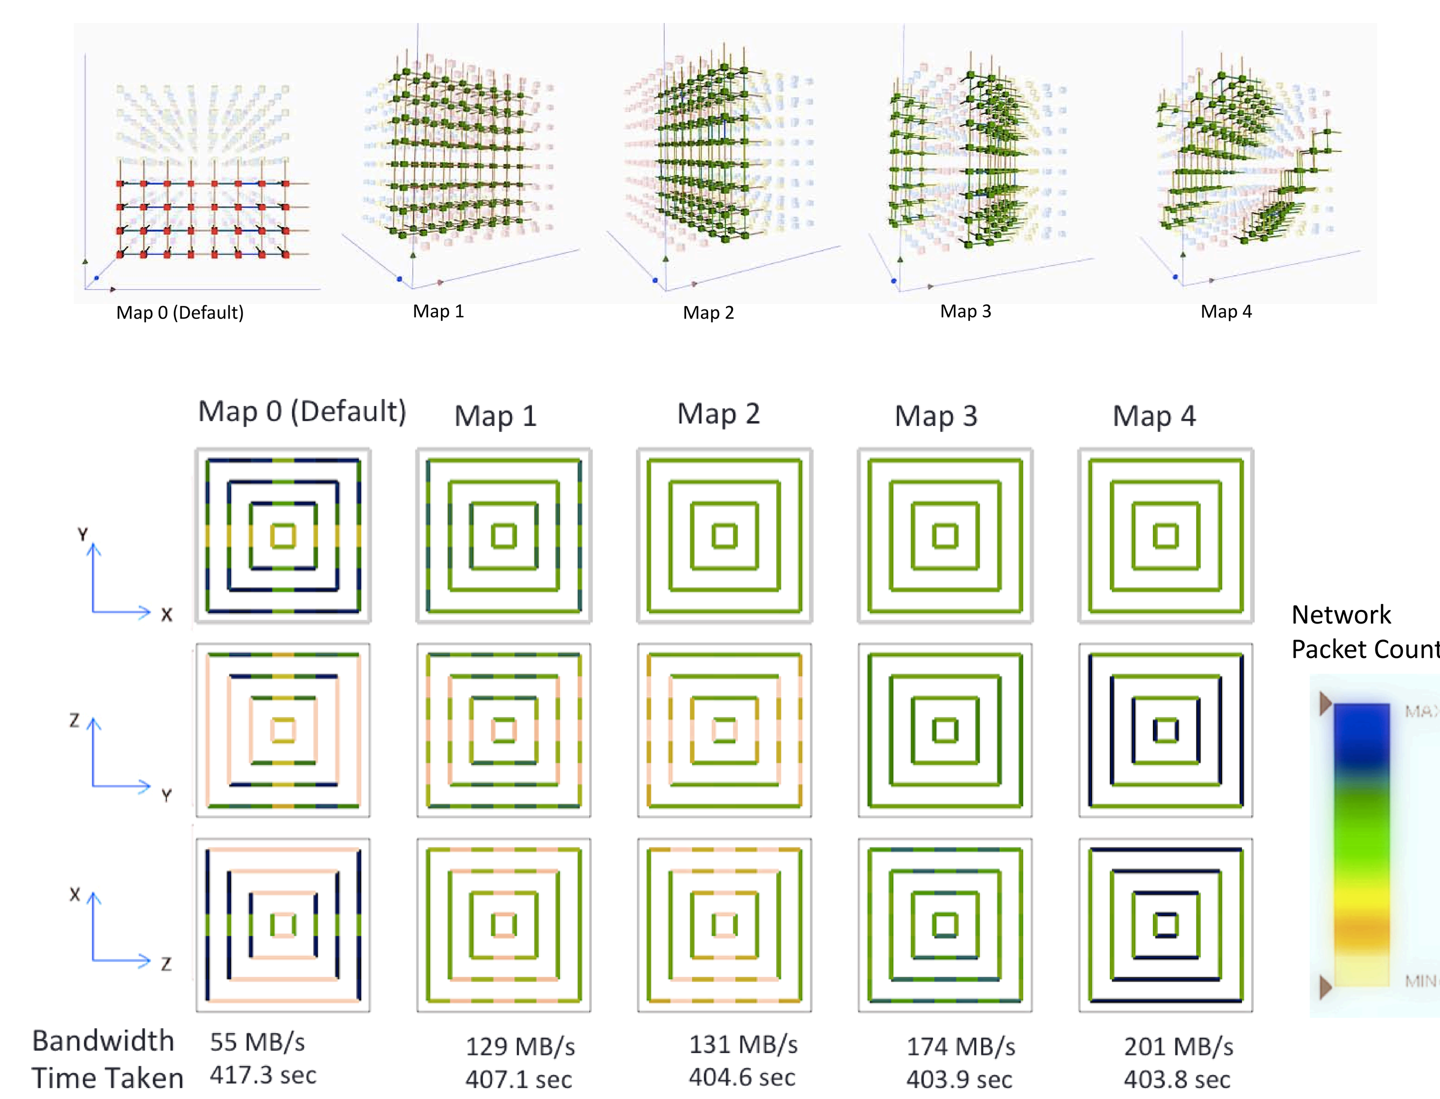

Communication patterns and network utilization for various node mappings for PF3D simulation. The various node maps (top) and the corresponding network traffic patterns shown in the aggregated views of the planar representation of the 3D torus network of IBM BG/P (bottom)

Projections of Hardware Domain onto the Communication Domain

Interesting aspects about the communication behavior of processes are revealed by projecting performance data on the virtual topology of applications

Scalable view of the SAMRAI timing data for 16K processes (right). The subtree with greater time values is expanded and shown with greater depth levels for analysis (left).

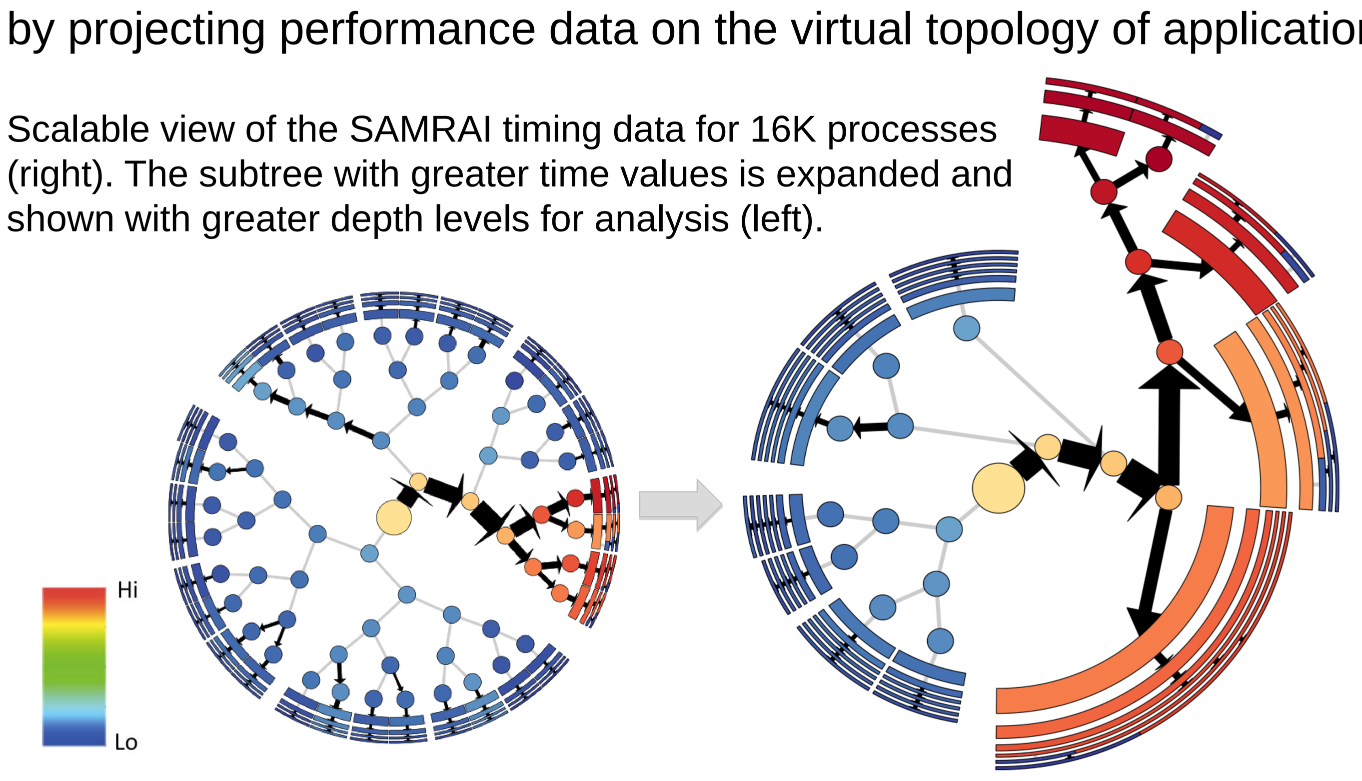

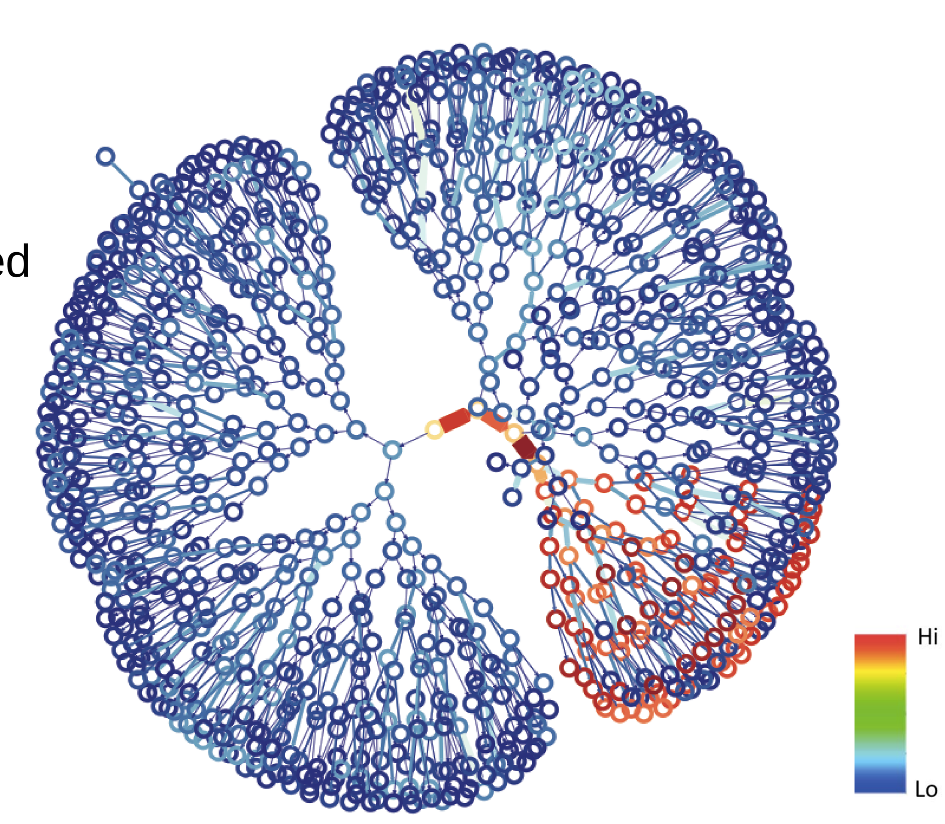

The load balancing timing data for SAMRAI projected on the communication graph. Every process is represented as a node and colored on the basis of time taken which clearly indicates the processes causing the delay (in red). Total 1024 processes.

References:

•M. Schulz et al. Interpreting performance data across intuitive domains. In *International Conference on Parallel Processing* (ICPP'11), Sept. 2011

•A. G. Landge et al. Visualizing network traffic to understand the performance of massively parallel simulations. *IEEE Transactions on Visualization and Computer Graphics*, 2012.

•A. Bhatele et al. Novel views of performance data to analyze large-scale adaptive applications. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, (SC '12) Nov. 2012.

•A. Bhatele et al. Mapping Applications with Collectives over Sub-Communicators on Torus Networks. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, (SC '12) Nov. 2012.

CEDMAV
CLEMSON School of COMPUTING
UC DAVIS UNIVERSITY OF CALIFORNIA
Lawrence Livermore National Laboratory
SCI www.sci.utah.edu
SCI X SCI INSTITUTE • EXHIBIT • EXPLORE • EXCITE • EXPERIENCE • EXCHANGE
THE UNIVERSITY OF UTAH