GCLab

SC19 Talks – Booth 575

Tuesday, 12:00 p.m.

Tellico HPC System: Unleashing High Performance Computing at UTK - Nigel Tan

Abstract: The University of Tennessee is the recipient of an IBM Shared University Research (SUR) Award that supported the acquisition of a high-performance computing cluster based on IBM Power9. The UTK cluster includes 4 compute 32-core Power9 nodes (256 GB RAM), two of which have 2 Nvidia V100 GPUs each. The cluster is supported by 1 elastic storage server (IBM Spectrum Scale); 2 TB RAW disk space; 1 Infiniband 36 port EDR TOR switch (non-blocking); and 1 IBM Ethernet switch (48x1Gb+4x10Gb) with 1 GB connectivity for the cluster. In this talk we give an overview of the system and show several use cases that are enabled by this Power9 system, including the simulation of the Gatlinburg wildfire using the Fire Dynamics Simulator (FDS) code and magnetic reconnection simulations using Vector Particle-In-Cell (VPIC) code.

Alternate time and location:

Thurs. 11-11:30 - IBM HPC & AI Univ. User Group Meeting Hyatt Regency/Centennial Ballroom F

Tuesday, 2:00 p.m.

Predicting Runtime and IO using Neural Networks
- Michael Wyatt

Abstract: For job allocation decisions, current batch schedulers have access to only information on the number of nodes and runtime. User-provided runtimes are typically inaccurate because users overestimate or lack understanding of job resource requirements. Beyond the number of nodes and runtime, other system resources, including IO and network, are not available but play a key role in system performance. We address the need for automatic, general, and scalable tools that provide accurate resource usage information to schedulers with PRIONN (Predicting Runtime and IO using Neural Networks). We automate prediction of per-job runtime and IO resource usage, enabling IO-aware scheduling and better resource management on HPC systems. The novelty of our tool is the input of whole job scripts into deep learning models, allowing complete automation of runtime and IO resource predictions. We demonstrate the power of PRIONN with runtime and IO resource predictions applied to IO-aware scheduling for real HPC data.

Tuesday, 3:00 p.m.

Performance Portable Plasma Simulations for the Exascale Era - Nigel Tan

Abstract: Vector Particle-In-Cell (VPIC) is a high performance plasma simulation code capable of scaling to the largest supercomputers in the world. VPIC has been optimized for CPU and many-core architectures but the recent diversifying of hardware magnifies the challenge. To keep up with the rapid growth in hardware, we ported VPIC using the Kokkos performance portability framework. This project quantifies performance and power tradeoffs that portability frameworks like Kokkos introduce. We use this information to better optimize VPIC application across multiple platforms including Power9. The initial port already achieves near ideal weak scaling on 2,048 Summit nodes with 12,288 GPUs.

Tuesday, 5:00 p.m.

SOMOSPIE: SOil MOisture SPatial Inference Engine
- Danny Rorabaugh

Abstract: The primary source of soil moisture data over large areas is satellite remote sensing technologies (i.e., radar-based systems), but such data have coarse resolution and often exhibit large spatial information gaps. Where data are too coarse or sparse for a given need (e.g., precision farming), one can leverage machine-learning techniques coupled with other sources of environmental information, such as topography or vegetation, to generate gap-free information at a finer spatial resolution. To this end, we develop a spatial inference engine consisting of modular stages for processing spatial environmental data, generating predictions with machine-learning techniques, and analyzing these predictions. We demonstrate the functionality of this approach and the effects of data processing choices via multiple prediction maps over a United States ecological region with a diverse soil moisture profile (i.e., the Middle Atlantic Coastal Plains).

THE UNIVERSITY OF TENNESSEE KNOXVILLE · XSEDE · Jetstream · Los Alamos NATIONAL LABORATORY EST. 1943 · IBM · NSF · GCL · UD · OAK RIDGE National Laboratory · ICL · USC

GCLab

SC19 Talks – Booth 575

Wednesday, 2:00 p.m.

Modeling Non-Determinism in HPC Applications - Dylan Chapp

Abstract: Runtime non-determinism in High Performance Computing (HPC) applications presents steep challenges for computational reproducibility and correctness. These challenges are magnified in the context of complex scientific codes where the links between observable non-determinism and root causes are unclear. We apply a three-phase workflow to (1) build graph-structured models of non-deterministic communication in parallel applications; (2) identify windows of execution with maximum run-to-run variability; and (3) map runtime non-determinism to source code level root causes.

Wednesday, 5:00 p.m.

Hermit: Reclaiming Lost Resources via Elastic Scheduling - Joe Teague

Abstract: Prior to resource reservations and the scheduling of large jobs, distributed computer systems using static schedulers undergo a drop in utilization for which a static scheduler cannot compensate. Complex workflows that utilize system resource reservations experience a ramp-up in resource usage over time. While the decrease in utilization could be used to get a head start on a workflow and reduce some of this ramp-up, static schedulers do not natively permit this type of allocation. To address this issue, we introduce Hermit: an extension to the Flux hierarchical scheduler that allows dynamic scheduling via growable resource allocations. Hermit is able to address both these problems by employing resources during drain events and allowing complex workflows to complete pre-processing before their reservation officially starts.

Thursday, 10:00 a.m.

Power Usage in Data-Intensive Applications using MapReduce over MPI - Paula Olaya

Abstract: Data analytics and data-intensive workloads are gaining representation at peta- and exascale. MapReduce has gained the most traction in the HPC community. **Mimir,** a novel **MapReduce over MPI** framework tackles skewed data, imbalance in memory usage, and loss in data scalability with combiner optimizations, dynamic repartitions, and a split method to handle datasets with superkeys. All this data movement **is power intensive** but little work is available in providing quantitative evaluations of these costs. This project quantitatively measures the impact of power capping on performance metrics such as runtime and power usage over time for data-intensive applications on top of a MapReduce over MPI framework when executed on HPC systems.

Thursday, 2:00 p.m.

Characterizing In Situ and In Transit Analytics of Molecular Dynamics Simulations for Next-generation Supercomputers - Michael Wyatt and Ian Lumsden

Abstract: Molecular Dynamics (MD) simulations executed on state-of-the-art supercomputers are producing data at rates faster than it can be written out to disk. In situ and in transit analysis of data generated by MD simulations reduce the original volume of information by several orders of magnitude, thereby alleviating the negative impact of I/O bottlenecks. This work focuses on characterizing the impact of in situ and in transit analytics on the overall MD workflow performance, and the capability for capturing rapid, rare events in the simulated molecular system. Our MD simulation and analysis processes share data using RDMA via DataSpaces. We measure the time spent waiting in I/O by the MD simulation and analysis to characterize the observed workflow patterns. The insights gained from this study are generally applicable for in situ and in transit workflows that require optimization of parameters to minimize loss in workflow performance and analytic accuracy.

THE UNIVERSITY OF TENNESSEE KNOXVILLE · XSEDE · Jetstream · Los Alamos NATIONAL LABORATORY EST. 1943 · IBM · NSF · GCL · OAK RIDGE National Laboratory · ICL · USC