ISSUE
24
SPRING
2017

QUARTERLY NEWSLETTER FOR THE COMPUTER SCIENCE AND MATHEMATICS DIVISION

CSMD

this issue

Science Highlights P.01
News P.13
Events P.13
Community Service P.18
Awards and Recognition P.19
Publications of Note P.19
About CSMD P.20

Science Highlights

Language-Based Optimizations for Persistence on Nonvolatile Main Memory Systems

Joel E. Denny, Seyong Lee, and Jeffrey S. Vetter

Researchers designed several novel language-based optimization techniques for programming NVM as persistent memory and demonstrated them as an extension of our NVL-C programming system.

This work enhances our ability to efficiently utilize NVM as high-performance, persistent memory in HPC systems.

Substantial advances in nonvolatile memory (NVM) technologies have motivated widespread integration of NVM into mobile, enterprise, and HPC systems. Recently, considerable research has focused on architectural integration of NVM and respective programming systems, exploiting NVM's trait of persistence correctly and efficiently. In this regard, we design several novel language-based optimization techniques for programming NVM and demonstrate them as an extension of our NVL-C system. Specifically, we focus on optimizing the performance of atomic updates to complex data structures residing in NVM. We build on two variants of automatic undo logging: canonical undo logging, and shadow updates. We show these techniques can be implemented transparently and efficiently, using dynamic selection and other logging optimizations. Our empirical results on several applications gathered on an NVM test bed illustrate that our cost-model-based dynamic selection technique can accurately choose the best logging variant across different NVM modes and input sizes. In comparison to statically choosing canonical undo logging, this improvement reduces execution time to as little as 53% for block-addressable NVM and 73% for emulated byte-addressable NVM on a Fusion-io ioScale device.

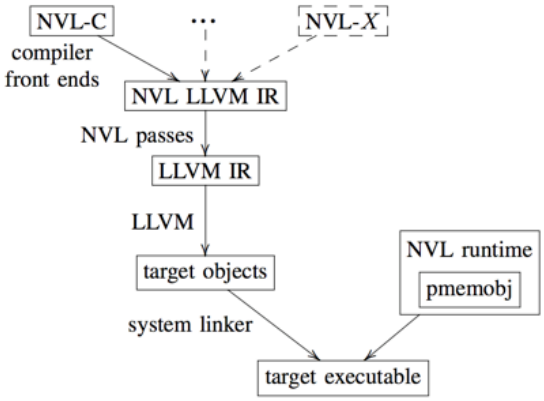

We extended NVL-C to support our techniques. As such, our compiler extensions are built on LLVM, encouraging reuse among compiler front ends, high-level languages, and system architectures. Our compiler extensions target Intel's pmemobj library for programming NVM.

This work was performed on ExCL in CSMD.

Publication: Joel E. Denny, Seyong Lee, and Jeffrey S. Vetter. "Language-Based Optimizations for Persistence on Nonvolatile Main Memory Systems." The 31st IEEE International Parallel and Distributed Processing Symposium (IPDPS 2017), to appear, Orlando, Florida, USA, May 2017. (accepted)

OAK RIDGE National Laboratory

Science Highlights (Continued)

Falcon: Visual Analysis of Large, Irregularly Sampled, and Multivariate Time Series Data in Additive Manufacturing

C. Steed, W. Halsey, R. Dehoff, S. Yoder, V. Paquit, and S. Powers

Researchers designed, developed, and deployed a new visual analytics system, called Falcon, that allows interactive exploration of large, time-series data using both data visualization and analytics.

Understanding and extracting information from lengthy and irregularly sampled time series presents a significant challenge in additive manufacturing as well as other domains. This work developed a new scalable tool, that provides researchers with visual overviews, the ability to explore trends, link observed behaviors and understand trends. In a case study with additive manufacturing experts, the use of Falcon resulted in new discoveries about the process at hand.

Flexible visual analysis of long, high-resolution, and irregularly sampled time series data from multiple sensor streams is a challenge in several domains. In the field of additive manufacturing, this capability is critical for realizing the full potential of large-scale 3D printers. This work proposes a visual analytics approach that helps additive manufacturing researchers acquire a deep understanding of patterns in log and imagery data collected by 3D printers. Specific goals include discovering patterns related to defects and system performance issues, optimizing build configurations to avoid defects, and increasing production efficiency. We introduce Falcon, a new visual analytics system that allows users to interactively explore large, time-oriented data sets from multiple linked perspectives. Falcon provides overviews, detailed views, and unique segmented time series visualizations, all with adjustable scale options. To illustrate the effectiveness of Falcon at providing thorough and efficient knowledge discovery, we present a practical case study involving experts in additive manufacturing and data from a large-scale 3D printer. Although the focus of this paper is on additive manufacturing, the techniques described are applicable to the analysis of any quantitative time series.

This work was performed ORNL and the ORNL MDF. Work sponsored by DOE/AMO.

Publication: C. Steed, W. Halsey, R. Dehoff, S. Yoder, V. Paquit and S. Powers, “Falcon: Visual analysis of large, irregularly sampled, and multivariate time series data in additive manufacturing,” Computers & Graphics, 63:50-64 (2017), http://dx.doi.org/10.1016/j.cag.2017.02.005

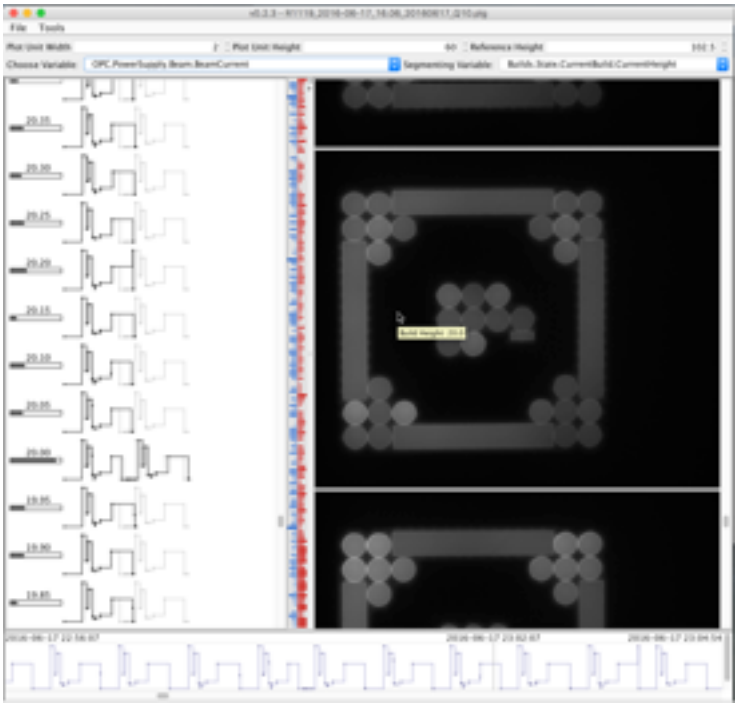

Falcon supports exploratory analysis of large and complex time series data using interactive data visualization and analytics. Here a detailed visualization highlights 3D printer layer build patterns using time and imagery information.

Science Highlights (Continued)

High-Performance Key-Value Store for HPC Systems

Neena Imam, Weikuan Yu, Huansong Fu, Ahana Roy Choudhury

Researchers developed Memcached for HPC Systems and provided solutions addressing the challenges to develop a key-value store for HPC systems.

This work demonstrates that an efficient key-value store can be developed for HPC systems by taking advantage of high-performing network and software abstractions available in the HPC systems.

There has been growing interest in enabling fast data analytics by leveraging system capabilities from large-scale high-performance computing (HPC) systems. OpenSHMEM is a popular programming model for HPC systems that has been used for large-scale compute-intensive scientific applications. In this work, we leverage OpenSHMEM to design a distributed in-memory key-value store for fast data analytics. To evaluate the idea, we develop SHMEMCache that takes advantage of OpenSHMEM's symmetric global memory, efficient one-sided communication operations and general portability. The results show that SHMEMCache outperforms in terms of latency and throughput over the original Memcached (shown in the graphs below). Our initial scalability test on the Titan supercomputer has also demonstrated that SHMEMCache can scale to 256 nodes and beyond.

This work was performed at Oak Ridge National Laboratory and Florida State University. It was sponsored by the US Department of Defense.

Publications: Huansong Fu, Manjunath Gorentla Venkata, Ahana Roy Choudhury, Neena Imam, Weikuan Yu, High-Performance Key-Value Store On OpenSHMEM, CCGRID, 2017.

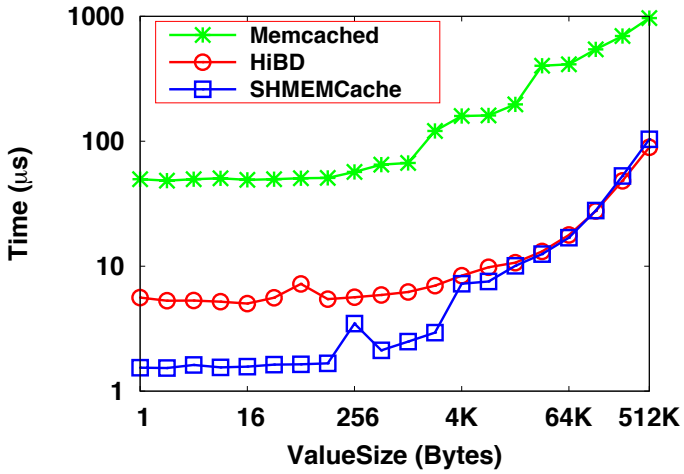

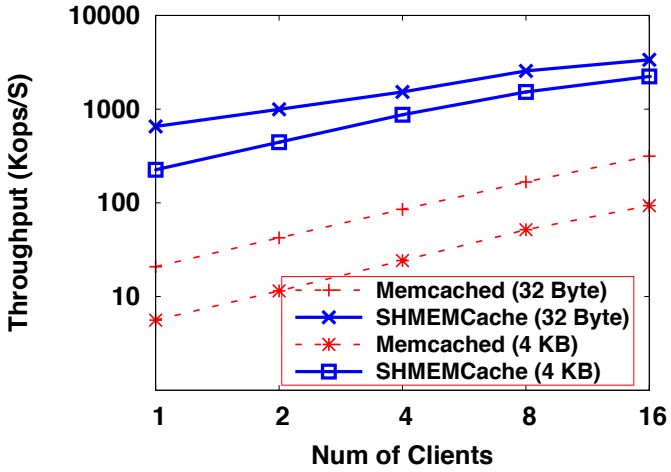

A Component-Level Power Profiler

Greg Watson, Greg Hinkel, and Chung-Hsing Hsu

Researchers designed and implemented a component-level power profiler.

This work enables a software developer to gain further insights about the power-consumption behavior of his/her codes because it provides component-level details. Having an energy-efficient code is important because energy efficiency is now a first-class design constraint in high performance computing.

Energy efficiency is a major challenge for high-performance computing (HPC), and the capability for accurate and fine-grained power monitoring is critical to the success of energy-efficient HPC. Solutions for mapping monitored power data back to the source code of a run are also important because they can help a software developer diagnose the inefficiency in the utilization of allocated resources. In this work we designed and implemented one such solution, a component-level power profiler. We essentially expand an industrial-strength, application-level performance profiler by incorporating component-level power data along with performance data collected by the profiler itself. In the process two bugs in the profiler were identified and in turn fixed by the vendor. We now have a capability for the power profiling of a parallel application at component level. Our initial evaluation indicates that the new profiler provides a reliable, higher sampling rate than Intel's Running Average Power Limit (RAPL) technology, a popular mechanism used in the HPC community. In addition, this profiler provides a complete picture of the power breakdown of a compute node, and this is something that RAPL is incapable of.

This work was supported by funding from the Department of Defense.

Presentation: Chung-Hsing Hsu, "Assessing the Quality of Component Level Power Profiling," Drumstrang Program, Jessup, MD. March 29, 2017.

Science Highlights (Continued)

Epidemic Failure Detection and Consensus for Extreme Parallelism

A. Katti, G. Di Fatta, T. Naughton, and C. Engelmann

Researchers developed three novel Gossip-based failure detection and consensus algorithms for User Level Failure Mitigation in Message Passing Interface applications to enable efficient handling of failures in extreme-scale high-performance computing systems.

This work significantly improves the performance of failure handing in extreme-scale high-performance computing systems, thus increasing the productivity of DOE's science applications.

Future extreme-scale high-performance computing systems will be required to work under frequent component failures. The Message Passing Interface (MPI) Forum's User Level Failure Mitigation (ULFM) proposal has introduced an operation, MPI_Comm_shrink, to synchronize the alive processes on the list of failed processes, so that applications can continue to execute even in the presence of failures by adopting algorithm-based fault tolerance techniques. This MPI_Comm_shrink operation requires a failure detection and consensus algorithm. We developed three novel failure detection and consensus algorithms using Gossiping.

Stochastic pinging is used to quickly detect failures during the execution of the algorithm, failures are then disseminated to all the fault free processes in the system and consensus on the failures is detected using the three consensus techniques. The algorithms were implemented and tested using the Extreme-scale Simulator. The results show that the stochastic pinging detects all the failures in the system. In all the algorithms, the number of Gossip cycles to achieve global consensus scales logarithmically with system size. The second algorithm also shows better scalability in terms of memory and network bandwidth usage and a perfect synchronization in achieving global consensus. The third approach is a three-phase distributed failure detection and consensus algorithm and provides consistency guarantees even in very large and extreme-scale systems while at the same time being memory and bandwidth efficient.

Work was performed at the University of Reading, UK and at ORNL. The work at ORNL was funded by DOE's Advanced Scientific Computing Research Office's Exascale Operating System and Runtime (ExaOS/R) program.

Publication: Amogh Katti, Giuseppe Di Fatta, Thomas Naughton, and Christian Engelmann, "Epidemic Failure Detection and Consensus for Extreme Parallelism," International Journal of High Performance Computing Applications (IJHPCA), February 2017. SAGE Publications. ISSN 1094-3420. DOI 10.1177/1094342017690910.

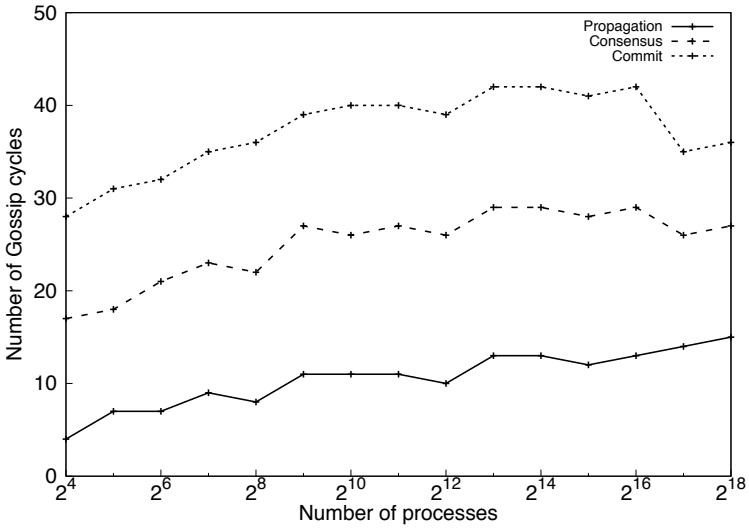

Number of Gossip cycles, i.e., the amount of time, needed for propagation, consensus and commit after a single failure injection (algorithm 3) with up to 262,144 participating processes.

Science Highlights (Continued)

Bridging Simulations with Neutron Scattering Yields New Capabilities to Probe Lipid Bilayers

J-M. Carrillo, J. Katsaras, B.G. Sumpter, and R. Ashkar

Researchers demonstrated a coarse-grained MD simulation approach that mimics neutron scattering data from large unilamellar lipid vesicles over a range of bilayer rigidities.

The ability to precisely probe the length and time scales that underlie ionic and molecular transport in lipid bilayers opens new opportunities to explore long-standing problems in membrane physics.

This work demonstrates a coarse-grained MD simulation approach that mimics neutron scattering data from large unilamellar lipid vesicles over a range of bilayer rigidities. While experimental observations of membrane structure and dynamics are limited in both spatial and temporal resolution, computer simulations when appropriately quantified/calibrated are able to intimately resolve dynamics and structure, enabling accurate interpretation of experimental results and a detailed molecular understanding of complex lipid bilayer membranes. In particular, the synergy between neutron scattering experiments and molecular simulations offers opportunities for new discoveries in membrane physics, as the length and time scales probed by molecular dynamics (MD) simulations parallel those of experiments. Specifically, in this work we have simulated vesicle form factors and membrane thickness fluctuations determined from small angle neutron scattering (SANS) and neutron spin echo (NSE) experiments, respectively. Our simulations accurately reproduce trends from experiments and thereby lay the groundwork for studies of more complex membrane systems.

This work was conducted at the Center for Nanophase Materials Sciences, which is a DOE Offi ce of Science User Facility. J. K. is supported by DOE' s Scientific User Facilities Division. R. A. is supported by ORNL' s Neutron Sciences Directorate Clifford G. Shull Fellowship. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

Publication: J-M. Carrillo, J. Katsaras, B.G. Sumpter, R. Ashkar. J. Chem. Theory & Comp, (2017). DOI: 10.1021/acs.jctc.6b00968.

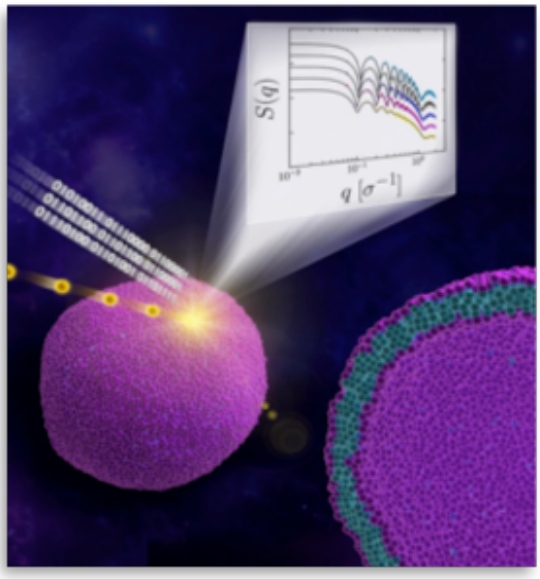

Schematic illustration of our new computational modeling capability for understanding dynamics of bilayer lipid membranes and for predicting neutron scattering results.

Science Highlights (Continued)

Cryogenic Memory Cell Design Based on Small Coupled Arrays of Josephson Junctions

Yehuda Braiman, Niketh Nair, Jake Rezac, Roland Glowinski, and Neena Imam

Researchers proposed cryogenic memory cell design that has a potential to substantially outperform the existing memory cells, achieve much faster access times and lower access and dissipation energies, and reduce the size of the memory cell.

Designing fast, energy efficient and small size memory circuit is considered one of the major bottlenecks in modern computing. Such memories could profoundly advance development and performance of exascale, quantum, and cryogenic computing.

We demonstrated a paradigm for cryogenic memory operation and presented a specific example of a circuit that consists of three inductively coupled Josephson junctions. We have employed Josephson junction parameter values that are consistent with the current state-of-the-art Josephson junction fabrication capabilities. For parameter values presented in the paper, memory cell access times are of the order of 10 - 100 ps while access energies are of the order of 0.1 – 5 aJ. The principles of memory cell design and operation described in our papers can in principle be implemented on other Josephson junction based circuits as well.

This work was supported by the United States Department of Defense and used resources from the Extreme Scale Systems Center, located at Oak Ridge National Laboratory. Oak Ridge National Laboratory is managed by UT- Battelle, LLC, under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (http://energy.gov/ downloads/doe-public-access-plan).

Publication: Y. Braiman, N. Nair, J. Resac, and N. Imam, “Memory Cell Operation Based on Small Josephson Junctions Arrays”, Superconductor Science and Technology 29, 124003 (2016); Y. Braiman, B. Neschke, N. Nair, N. Imam, and R. Glowinski, “Memory States in Small Arrays of Josephson Junctions” Physical Review E 94, 052223 (2016); J. D. Rezac, N. Imam, and Y. Braiman, “Parameter Optimization for Transitions between Memory States in Small Arrays of Josephson Junctions”, Physica A 474, 267 (2017).

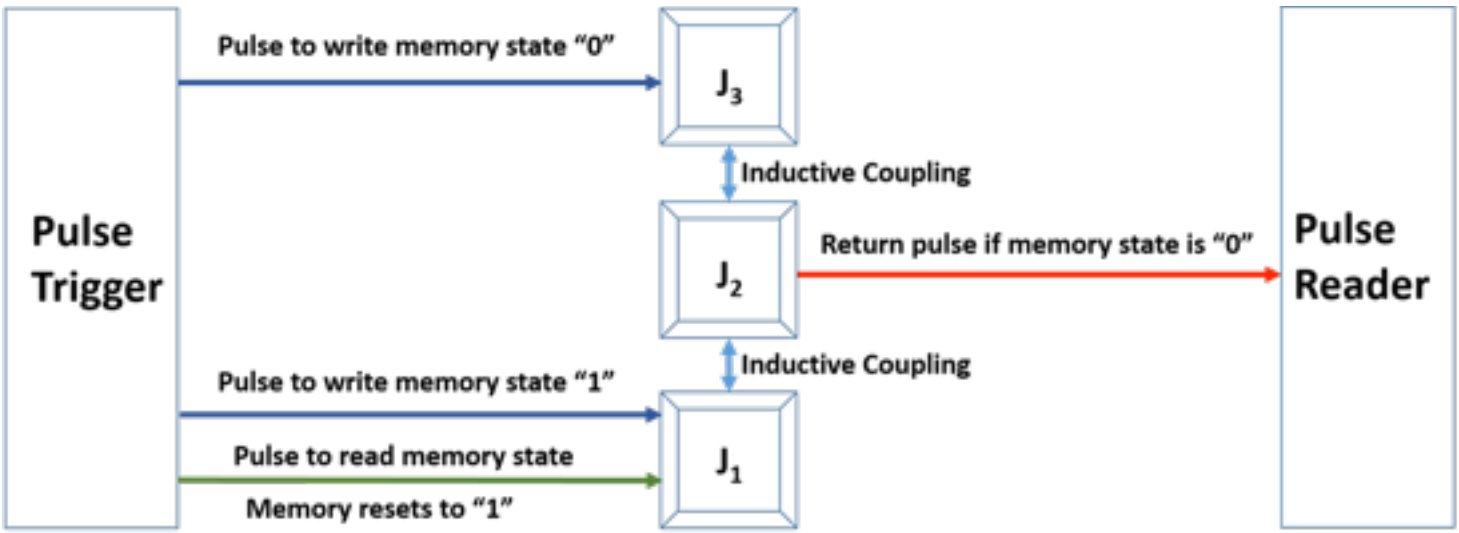

The principles of memory operation of the schematic design presented in Figure are as follows: an appropriate pulse sent to the junction J1 will result in writing memory state ‘1’ if the previous memory state was ‘0’ and will not change the memory state if the previous memory state was ‘1’. A suitable pulse sent to the junction J3 will result in writing memory state ‘0’ (or memory reset) if the previous memory state was ‘1’ and will not change the memory state if the previous memory state was ‘0’. Reading the memory state will be implemented by sending a proper pulse to the junction J1. If the memory state was ‘0’, junction J2 will respond by generating a large amplitude voltage pulse that can be read by pulse reader and further processed. If, however, the memory state was ‘1’ no noticeable voltage pulse will be generated from the junction J2.

Science Highlights (Continued)

Design and Implementation of Papyrus: Parallel Aggregate Persistent Storage

Jungwon Kim, Seyong Lee, Joel E. Denny, and Jeffery S. Vetter

Researchers developed a novel programming system for aggregating and using distributed NVM within an HPC system.

This work enhances our ability to effectively utilize emerging NVM technologies, which the upcoming DOE systems will all have.

A surprising development in recently announced HPC platforms is the addition of, sometimes massive amounts of, persistent (nonvolatile) memory (NVM) in order to increase memory capacity and compensate for plateauing I/O capabilities. However, there are no portable and scalable programming interfaces using aggregate NVM effectively. This paper introduces Papyrus: a new software system built to exploit emerging capability of NVM in HPC architectures. Papyrus (or Parallel Aggregate Persistent -YRU- Storage) is a novel programming system that provides features for scalable, aggregate, persistent memory in an extreme-scale system for typical HPC usage scenarios. Papyrus mainly consists of Papyrus Virtual File System (VFS) and Papyrus Template Container Library (TCL). Papyrus VFS provides a uniform aggregate NVM storage image across diverse NVM architectures. It enables Papyrus TCL to provide a portable and scalable high-level container programming interface whose data elements are distributed across multiple NVM nodes without requiring the user to handle complex communication, synchronization, replication, and consistency model. We evaluate Papyrus on two HPC systems, including UTK Beacon and NERSC Cori, using real NVM storage devices.

This work was performed UTK Beacon and NERSC Cori Phase I.

Publication: Jungwon Kim, Kittisak Sajjapongse, Seyong Lee, and Jeffrey S. Vetter. “Design and Implementation of Papyrus: Parallel Aggregate Persistent Storage”. The 31st IEEE International Parallel and Distributed Processing Symposium (IPDPS 2017), to appear, Orlando, Florida, USA, May 2017. (accepted)

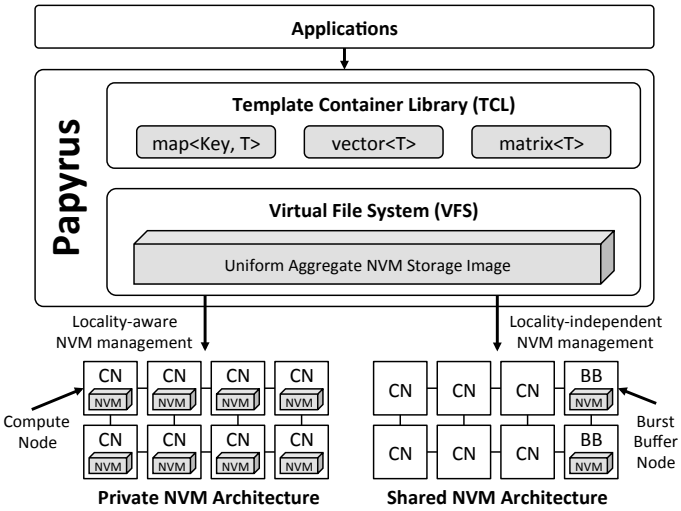

Papyrus consists of Papyrus Virtual File System (VFS) and Template Container Library (TCL). It runs on both of private NVM architecture (i.e., ORNL’s Summit, LLNL’s Sierra) and shared NVM architecture (i.e., NERSC’s Cori and LANL/ Sandia’s Trinity).

Science Highlights (Continued)

Exploring the Effect of Compiler Optimizations on the Reliability of HPC Applications

R. A. Ashraf, R. Gioiosa (Pacific Northwest National Laboratory - PNNL), G. Kestor (PNNL), and R. F. DeMara (University of Central Florida)

Researchers demonstrated that compiler optimizations have an impact on the vulnerability of HPC applications in addition to performance, and therefore the need to consider tradeoffs between performance and vulnerability for critical application components.

The strict power efficiency required to achieve exascale systems will result in an increase in the number of detected and undetected transient errors, which would directly affect the vulnerability of High-Performance Computing applications executed on these systems. Among the various hardware and software components that affect system resiliency, the impact of compiler optimizations on application vulnerability has not been widely investigated. This is significant since compiler optimizations are one of the most successful ways to improve performance with relatively low effort from the programmer. We analyze the tradeoffs between performance and vulnerability for multiple HPC applications compiled using standardized optimization levels. Results for these applications demonstrate that highly-optimized code is generally more vulnerable than un-optimized code. Increasing optimization levels can drastically improve performance as expected, however, it is observed that certain cases of optimization only provide marginal benefits in application performance yet considerably increase application vulnerability. In this work, we also define software attributes which are effected by changing compiler optimization levels, and contribute to increase in application vulnerability. For example, one of the attribute captures the probability of a faulty value to propagate to memory and corrupt the memory state of other processes in a parallel environment. Overall, our work shows that compiler cost functions should account for vulnerability in addition to standard metrics of code size and performance, and the importance of identifying important software attributes for this purpose.

This work was performed partially at Oak Ridge National Laboratory (ORNL) and was sponsored in part by the ORNL Postdoctoral Professional Development Program.

Publication: R. A. Ashraf, R. Gioiosa, G. Kestor, and R. F. DeMara, “Exploring the Effect of Compiler Optimizations on the Reliability of HPC Applications,” 22nd IEEE Workshop on Dependable Parallel, Distributed and Network-Centric Systems, May 29-June 2, 2017, Orlando, FL (PTS # 74150).

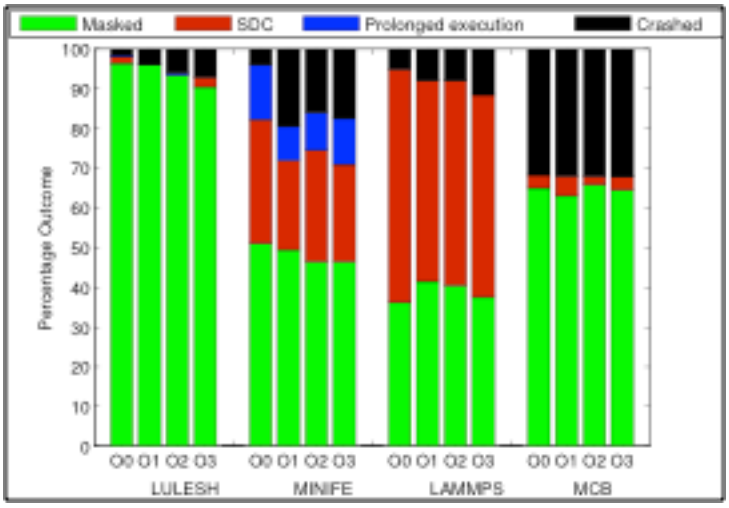

The impact of soft error injection on the outcome of multiple HPC applications (LULESH, miniFE, LAMMPS, and MCB) compiled using different optimization levels. Higher crash rate is observed at higher optimization levels for most applications.

Science Highlights (Continued)

Towards the Adaptation of the Graph500 benchmark to the SHARed data-structure-centric Programming (SharP) paradigm

Ferrol Aderholdt, Manjunath Gorentla Venkata, and Zachary W. Parchman – CSMD, Oak Ridge National Laboratory

Researchers designed, implemented, and preliminarily evaluated the adaptation of the Graph500 benchmark to SharP with a data-centric focus.

The adaptation of the Graph500 benchmark shows the significance of data-locality with respect to data placement and the capabilities of SharP.

The Graph500 benchmark is a data-intensive benchmark meant to be a compliment to the compute-intensive benchmark commonly used for ranking supercomputers within the Top500 list (i.e, the High Performance Linpack (HPL) benchmark). The Graph500 benchmark constructs a large undirected graph and performs multiple breadth-first search (BFS) operations on this graph. After the BFS operations are completed, the benchmark proceeds to validate the spanning tree that was created to ensure proper traversals. With respect to data usage, the benchmark provides low spatial and temporal locality and views memory as flat with no support for the heterogeneous and hierarchical memory that will be present in future extreme-scale systems. Additionally, the current design and implementation of the MPI-2 one-sided reference implementation does not determine the locality of the data to the nearest high performance NIC, which will provide higher latency for operations on data items far from the NIC when compared to those close to the NIC. This is important as the benchmark itself is extremely latency sensitive. This work adapts the MPI-2 one-sided reference implementation to leverage the SharP library, which allows for memory to be allocated near the NIC and structures data on behalf of the user. The implementation of this approach was preliminarily evaluated on the Titan supercomputer located at the OLCF comparing the benchmark with its MPI-2 one-sided reference implementation and an implementation in Cray SHMEM. We found the usage of SharP and focusing on a data-centric approach increased the performance of the benchmark with respect to traversed edges per second (TEPS) by as much as 11x with respect to Cray SHMEM and roughly 18x when compared to MPI.

This work was performed with support from ORNL.

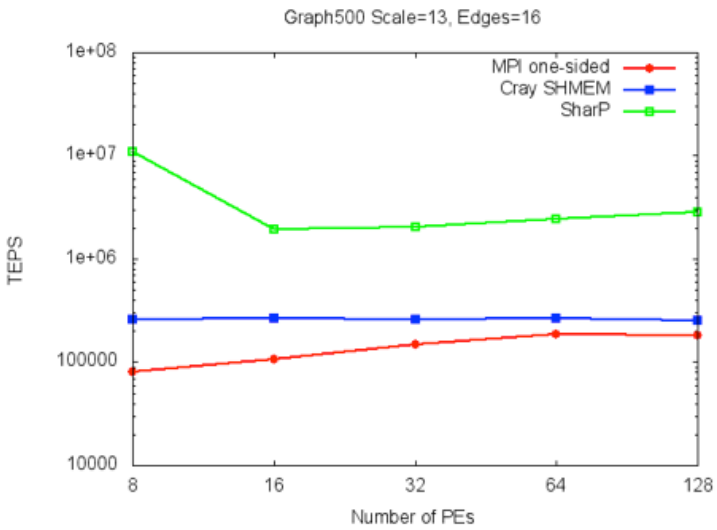

Science Highlights (Continued)

Time Synchronization in the Electric Power System

S. Achanta, M. Danielson, P. Evans, T. Jones, H. Kirkman, Y. Li-Baboud, R. Orndorff, A. Silverstein, K. Thomas, G. Trevino, F. Tuffner, and M. Weiss

We identified and articulated what power system engineers and operators need to know about the role and emerging importance of high-quality timing sources in routine and mission-critical grid applications. We articulated architectures and guidelines for how to get timing right and adopt best practices for timing, systems, and equipment that provide the correct, consistent time needed for emerging mission-critical applications, such as closed loop wide-area controls and real-time operator decision support.

PMUs need access to reliable Coordinated Universal Time (UTC) to allow synchrophasor applications to time-align the voltage and current time series data for analysis and coordinated activity over a wide geographical area. Most of the PMUs in North America acquire timing data from GPS. Since GPS signals are weak, they are inherently vulnerable to both human-caused and natural disruptions.

The power system uses precision timing for grid monitoring and situational awareness, to coordinate the operation and integration of a variety of grid assets, and for grid protection and operation. Because power systems are so large and often geographically separated, power data acquisition systems need to share a common time source. In some advanced applications, clocks need to be precisely synchronized to a common reference to enable the integration of diverse data types and sources and assure that decentralized, parallelized analysis and control actions are effectively coordinated and implemented. Information about time is used to understand how relationships and conditions change over time, to identify and understand specific events (whether an individual lightning strike or the sequence of events for a system collapse), and to understand dynamic and transient grid events.

This work was a collaboration between ORNL, other laboratories, industry and academia. It was sponsored by the US Department of Energy Office of Electricity Delivery.

Publication: S. Achanta, M. Danielson, P. Evans, T. Jones, H. Kirkman, Y. Li-Baboud, R. Orndorff, A. Silverstein, K. Thomas, G. Trevino, F. Tuffner, and M. Weiss, "Time Synchronization in the Electric Power System", North American synchrophasor Initiative, Tech Report NASPI-2017-TR-001, Richland, WA, March 2017.

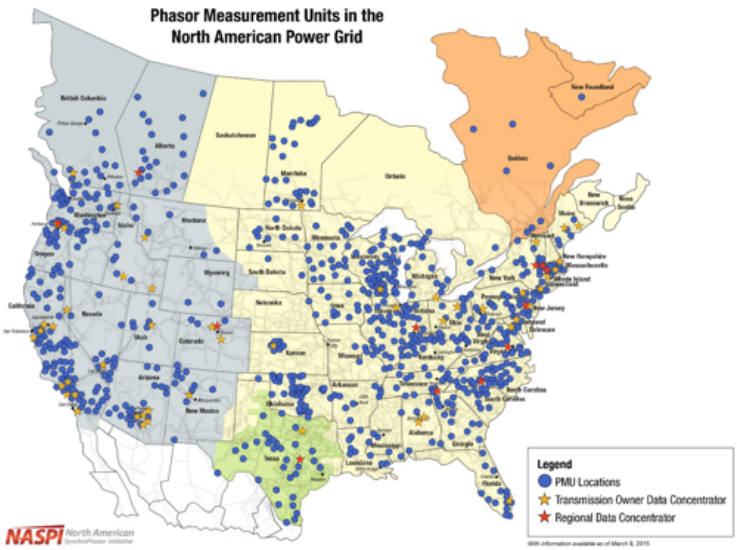

Almost 1,800 PMUs were installed across North America in 2015.

Science Highlights (Continued)

A New Ice Sheet Model Validation Framework for Greenland

Price, S.F., M.J. Hoffman, J.A. Bonin, I.M. Howat, T. Neumann, J. Saba, I. Tezaur, J. Guerber, D.P. Chambers, K.J. Evans, J.H. Kennedy, J. Lenaerts, W.H. Lipscomb, M. Perego, A.G. Salinger, R.S. Tuminaro, M.R. van den Broeke, and S.M.J. Nowicki

Researchers bridge the expertise gap between modelers and observationalists by using the Cryospheric Model Comparison Tool (CmCt) to evaluate satellite-based observations of polar ice sheet change.

We propose a new ice sheet model validation framework – the CmCt – that takes advantage of ice sheet altimetry and gravimetry observations collected over the past several decades and is applied here to modeling of the Greenland ice sheet. We use realistic simulations performed with the Community Ice Sheet Model (CISM) along with two idealized, non-dynamic models to demonstrate the framework and its use. Dynamic simulations with CISM are forced from 1991 to 2013, using combinations of reanalysis-based surface mass balance and observations of outlet glacier flux change. We propose and demonstrate qualitative and quantitative metrics for use in evaluating the different model simulations against the observations. We find that the altimetry observations used here are largely ambiguous in terms of their ability to distinguish one simulation from another. Based on basin-scale and whole-ice-sheet-scale metrics, we find that simulations using both idealized conceptual models and dynamic, numerical models provide an equally reasonable representation of the ice sheet surface (mean elevation differences of < 1 m). This is likely due to their short period of record, biases inherent to digital elevation models used for model initial conditions, and biases resulting from firn dynamics, which are not explicitly accounted for in the models or observations. On the other hand, we find that the gravimetry observations used here are able to unambiguously distinguish between simulations of varying complexity, and along with the CmCt, can provide a quantitative score for assessing a particular model and/or simulation. The new framework demonstrates that our proposed metrics can distinguish relatively better from relatively worse simulations and that dynamic ice sheet models, when appropriately initialized and forced with the right boundary conditions, demonstrate a predictive skill with respect to observed dynamic changes that have occurred on Greenland over the past few decades. An extensible design will allow for continued use of the CmCt as future altimetry, gravimetry, and other remotely sensed data become available for use in ice sheet model validation.

Support for Stephen F. Price, Matthew J. Hoffman, Irina Tezaur, Katherine J. Evans, Joseph H. Kennedy, William H. Lipscomb, Mauro Perego, Andrew G. Salinger, and Raymond S. Tuminaro was provided through the Scientific Discovery through Advanced Computing (SciDAC) program funded by the US Department of Energy (DOE), Office of Science, Advanced Scientific Computing Research and Biological and Environmental Research Programs. Support for Ian M. Howat was provided by NASA Cryospheric Sciences grant NNX11AR47G. Stephen F. Price was also partially supported by NASA Cryospheric Sciences, and Stephen F. Price and Matthew J. Hoffman were partially supported by the National Science Foundation, under grant ANT-0424589 to the Center for Remote Sensing of Ice Sheets (CReSIS). Jan Lenaerts and Michiel R. van den Broeke acknowledge funding from the Polar Program of the Netherlands Organization for Scientific Research (NWO/NPP), the Netherlands Earth System Science Center (NESSC) and Water, Climate, Ecosystems research theme of Utrecht University. Jan Lenaerts is supported by NWO ALW through a Veni postdoctoral grant. This research used resources of the National Energy Research Scientific Computing Center (NERSC; supported by the Office of Science of the US Department of Energy under contract DE-AC02-05CH11231).

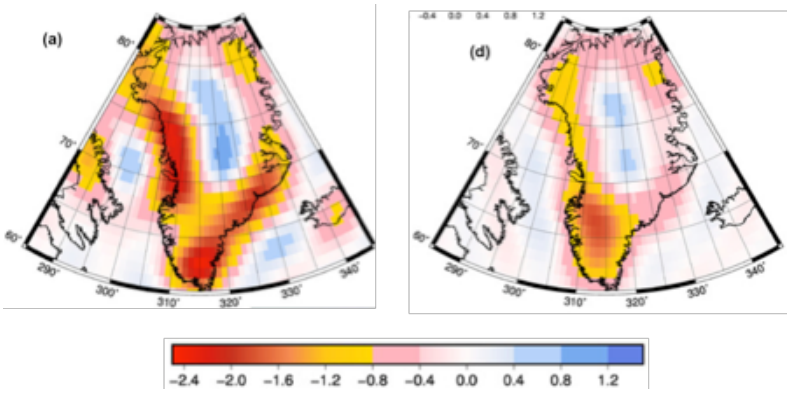

Map of Greenland mass trend between 2003-2013 (in m of water equiv. height) from GRACE observations (left) and from a high-resolution, CISM-Albany simulation (right).

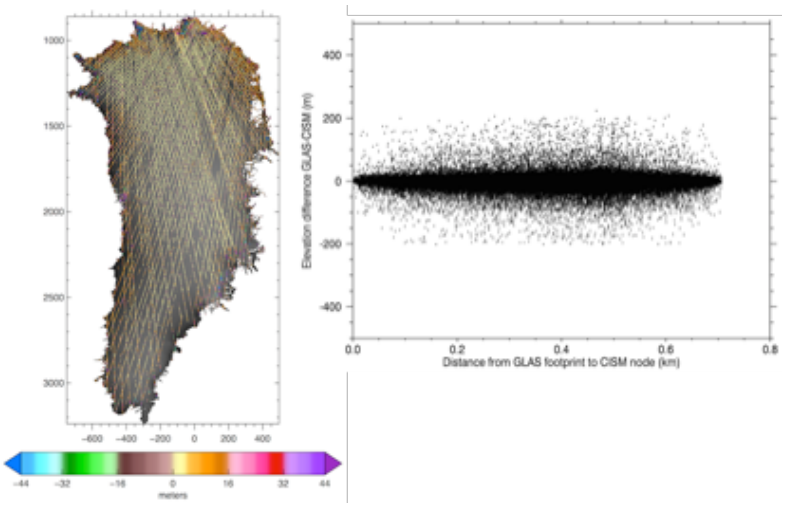

Map of ICESat vs. model elevation differences in 2007 (left) and scatter plot of the same (right).

Publication: Price, S.F., M.J. Hoffman, J.A. Bonin, I.M. Howat, T. Neumann, J. Saba, I. Tezaur, J. Guerber, D.P. Chambers, K.J. Evans, J.H. Kennedy, J. Lenaerts, W.H. Lipscomb, M. Perego, A.G. Salinger, R.S. Tuminaro, M.R. van den Broeke, and S.M.J. Nowicki. 2017. An ice sheet model validation framework for the Greenland ice sheet. Geophysical Model Development, 10, doi:10.5194/gmd-10-255-2017 .

Science Highlights (Continued)

GPU Acceleration of History-Based Multigroup Monte Carlo

Steven Hamilton (RNSD), Thomas Evans (RNSD), Stuart Slattery (CSMD)

Researchers developed scalable threaded algorithms for Monte Carlo neutron transport on the GPU. Work awarded Mathematics & Computation Division Best Summary+Presentation Award at ANS 2016 Winter Meeting.

Understanding how traditional Monte Carlo neutron transport algorithms map to multiple GPU architectures and how algorithmic variants can improve performance.

Accurate solutions of the radiation transport equation play an important role in the analysis of many nuclear systems. Because there is no need to discretize the problem, Monte Carlo methods offer the most accurate possible solutions to the trans- port equation. However, significant computational effort is required to reduce the stochastic noise inherent to Monte Carlo methods. Current trends in high performance computing are moving towards vectorized, single instruction multiple data (SIMD), architectures such as Intel Xeon Phis and graphics processing units (GPUs) due to their favorable ratio of performance to power consumption. The challenge of adapting Monte Carlo transport to vectorized computing architectures was first addressed by Brown with an algorithm known as event-based Monte Carlo. In this approach, rather than simulating individual particle histories from birth to death, groups of particles are processed one event at a time. This allows each event batch to perform the same operations in lockstep, enabling use of the vector hardware. Although modern vectorized hardware is more versatile than the early vector machines, reducing so-called thread divergence is nonetheless still an important consideration for algorithm design. Many of the algorithms proposed so far for Monte Carlo transport on advanced architectures have focused on variants of event-based transport. Some evidence, however, suggests that low- level thread divergence, resulting from short-lived branching statements, may not be detrimental to overall performance on GPUs if the effect on achievable memory bandwidth is considered. This paper characterizes the performance of a traditional history-based Monte Carlo algorithm on GPUs. We show that with only minor modifications to traditional algorithms, significant performance gains relative to CPUs are possible.

This work was supported by the ORNL LDRD program. This manuscript has been authored by UT- Battelle, LLC, under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (http://energy.gov/ downloads/doe-public-access-plan).

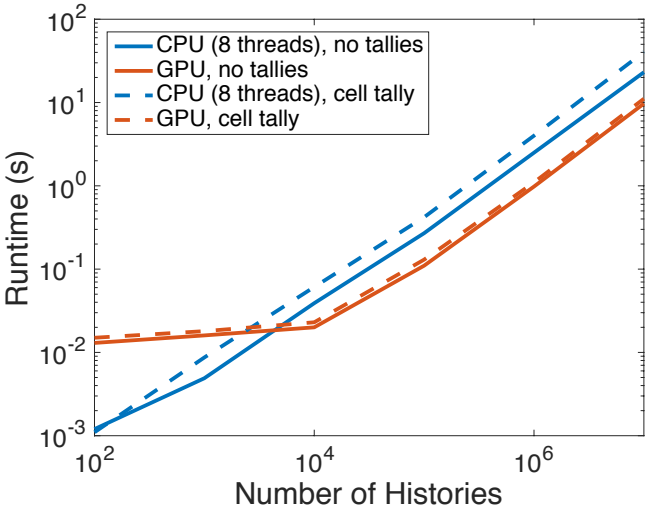

Comparison of CPU and GPU runtimes with and without tallies enabled for an increasing number of neutrons in the simulation. Each thread computes the trajectory of a neutron from birth to death as it interacts with the surrounding material in a nuclear reactor. Tallies allow for the measurement of neutron flux and other physics-based responses over the course of the simulation. It was found that low-level thread divergence was not as detrimental to GPU performance as previously thought. In addition, the atomic operations needed to perform the running tally sums over the course of the transport calculation were found to not substantially impact runtimes in the GPU calculation.

Publication: "GPU Acceleration of History-Based Multigroup Monte Carlo"; Steven Hamilton, Thomas Evans, Stuart Slattery; 2016 ANS Winter Meeting (2016).

News

Small nanoparticles have surprisingly big effects on polymer nanocomposites

Polymer nanocomposites mix particles billionths of a meter (nanometers, nm) in diameter with polymers, which are long molecular chains. Often used to make injection-molded products, they are common in automobiles, fire retardants, packaging materials, drug-delivery systems, medical devices, coatings, adhesives, sensors, membranes and consumer goods. When a team led by the Department of Energy's Oak Ridge National Laboratory tried to verify that shrinking the nanoparticle size would adversely affect the mechanical properties of polymer nanocomposites, they got a big surprise.

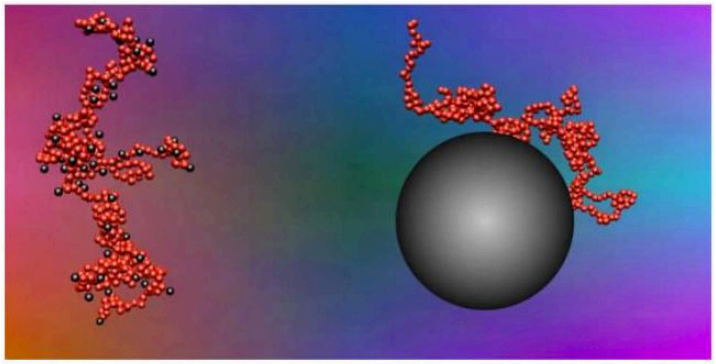

Small nanoparticles stick to segments of polymer chain that are about the same size as the nanoparticles themselves; these interactions produce a polymer nanocomposite that is easier to process because nanoparticles move fast, quickly making the material less viscous. At right, many segments of a polymer chain stick to a larger nanoparticle, making it difficult for that nanoparticle to move. Its slower movement results in a viscous material that is more difficult to process. Image credit - Oak Ridge National Laboratory, U.S. Dept. of Energy, created by Jan-Michael Carrillo

Read the full story here - https://www.ornl.gov/news/small-nanoparticles-have-surprisingly-big-effects-polymer-nanocomposites.

Diamonds that deliver

Neutrons, simulation analysis of tRNA-nanodiamond combo could transform drug delivery design principles

For drugs to be effective, they have to be delivered safely and intact to affected areas of the body. And drug delivery, much like drug design, is an immensely complex task. Cutting-edge research and development like that conducted at the U.S. Department of Energy's Oak Ridge National Laboratory can help solve some of the challenges associated with drug delivery.

In fact, ORNL researchers and collaborators at Wayne State University recently used a unique combination of experimentation and simulation to shed light on the design principles for improved delivery of RNA drugs, which are promising candidates in the treatment of a number of medical conditions including cancers and genetic disorders. Specifically, the research team discovered that the motions of a tRNA (or transfer RNA) model system can be enhanced when coupled with nanodiamonds, or diamond nanoparticles approximately 5 to 10 nanometers in size.

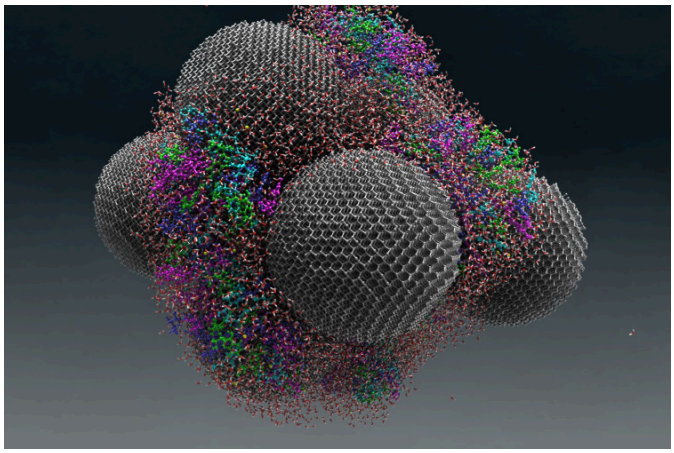

Read the full story here - https://www.ornl.gov/news/diamonds-deliver

Events

2nd International Workshop on Performance Portable Programming Models for Accelerators (P^3MA)

The workshop provides a forum bringing together researchers and developers to examine heterogeneous computing and how it has been evolving across an increasingly diverse set of accelerated architectures. Including an invited opening keynote address and a closing Q&A panel with all presenters, this workshop will provide perspectives from current research and a chance for attendees to actively participate in this quickly changing and growing area of HPC research.

This workshop will be held on June 22, 2017, colocated with the ISC High Performance Conference in Frankfurt, Germany (http://www.isc-hpc.com).

CAM Seminar Series

Dr. George Fann: Multiresolution Representations and Nonlinear Approximations of Operators with Boundary Conditions

We develop a multiresolution representation of a class of integral operators satisfying boundary conditions in order to construct fast O(N)-type algorithms for their applications. Some theoretical issues related to the construction of periodic Green's function for the Poisson equation will be described.

We illustrate the results on several examples in one and

three dimensions: the Hilbert transform, the non-oscillatory Helmholtz Green's function and the Poisson operator.

Dr. Jeremy Trageser: Nonlocal Models and Their Applications to the Study of Materials

This talk will explore nonlocal models and a few of their applications, particularly with regard to damage propagation and nucleation in materials. Through the nonlocal theory of peridynamics, damage within various material can be simulated quite naturally due to the lack of requirement of spatial differentiability in displacement fields. We will discuss a general peridynamic model which can be utilized to model fully anisotropic materials, covering any of the material symmetries in classical linear elasticity. This model allows us to model and analyze several new classes of materials through the lens of peridynamics.

We will also consider a nonlocal fourth-order model of plate vibration and discuss some of its connections with its local counterpart. This will include the introduction of two common types of boundary conditions, hinged and clamped, that may be imposed on the nonlocal problem. We will further strengthen the connection of the nonlocal boundary-valued problems with their local counterparts by showing the solutions of the nonlocal problems converge to the solutions of the local problems when the nonlocal interaction goes to zero.

Dr. Lukas Vlcek: Consistent Integration of Experimental and ab initio Data into Effective Physical Models

As the study of natural processes and discovery of new materials increasingly relies on computational approaches, the accuracy of the underlying models grows in importance. Some of the main obstacles in this pursuit are the missing agreed upon criteria for the evaluation of model quality and the lack of a rigorous basis for consistent integration of experimental and simulation data. To address these issues, we have recently developed general principles for the approximation of statistical mechanical systems based on the idea of distinguishability between the results of physical measurements, which naturally incorporate the resolution and sampling characteristics of the data. The corresponding metric quantifying distinguishability between a model and its target system can be then used as a robust basis for the construction of arbitrary statistical mechanical models, and as an objective measure of their quality. Implementation of the methodology takes advantage of an efficient perturbation technique for scanning model parameter spaces. In this presentation I will discuss our approach to consistent integration of experimental and ab initio data from diverse sources on the examples of atomistic force fields for thermodynamic systems and coarse-grained models used to extract the physical driving forces behind structures observed in microscopic imaging data.

Dr. Hoang Tran: Improving Recovery Guarantees for Sparse Approximations with Convex and Nonconvex Regularizations

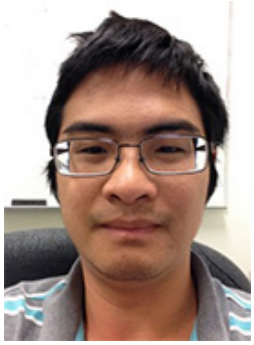

In this talk, we present some new recovery guarantees for sparse approximations via several types of convex and nonconvex optimizations. First, we show an improved sufficient condition for sparse Legendre expansions, which does not require the uniform bound of the underlying basis. Our sample complexity, independent of maximum polynomial degree, is established using the restricted eigenvalue property and the unbounded envelope of all Legendre polynomials. The analysis reveals some easy-to-test criteria for random sample sets under which the reconstruction error can be improved. Secondly, we present a new result for compressed sensing showing that nonconvex minimizations are at least as good as l1 minimization in exact recovery of sparse signals. This fact, while observed in many numerical experiments and proved for a few specific cases, is not completely verified in theory. We develop a recovery guarantee through a unified null space property that encompasses most currently proposed nonconvex functionals in literature. Several nonconvex functionals will be explored and the specific conditions in order to guarantee the recovery will be given.

Dr. Lili Ju: Stabilized Compact Exponential Time Differencing Methods for Gradient Flows and Scalable Implementation

In this talk, we present stabilized compact exponential time differencing (cETD) methods for numerical solutions of a family of gradient flow problems, which have wide applications in materials science, fluids and biology researches. These problems often form a special class of parabolic equations of different orders with high nonlinearity and stiffness, thus are very hard to solve efficiently and robustly over large space and time scales. The proposed methods achieve efficiency, accuracy and provable energy stability under large time stepping by combining linear operator splittings, compact discretizations of spatial operators, exponential time integrators, multistep or Runge-Kutta approximations and fast Fourier transform. We also develop a scalable

Events (Continued)

implementation of the cETD methods based on domain decomposition and operator localization with appropriate boundary conditions for subdomains, which is highly suitable for parallel computing. Various numerical experiments are carried out to demonstrate superior performance of the methods, including extreme scale phase field simulations of coarsening dynamics on the Sunway TaihuLight supercomputer.

Dr. M. Paul Laiu: Positivity Limiters for Filtered Spectral Approximations of Linear Kinetic Transport Equations

We analyze the properties and compare the performance of several positivity limiters for spectral approximations with respect to the angular variable of linear transport equations. It is well-known that spectral methods suffer from the occurrence of (unphysical) negative spatial particle concentrations due to the fact that the underlying polynomial approximations are not always positive at the kinetic level. Positivity limiters address this defect by enforcing positivity of the polynomial approximation on a finite set of preselected points. With a proper PDE solver, they ensure the positivity of the particle concentration at each step in a time integration scheme. We review several known positivity limiters proposed in other contexts and also introduce a modification for one of them. We give error estimates for the consistency of the positive approximations produced by these limiters and compare the theoretical estimates to numerical results. We then solve two benchmark problems with these limiters, make qualitative and quantitative observations about the solutions, and then compare the efficiency of the different limiters.

Mr. Armenak Petrosyan: Dynamical Sampling and Systems from Iterations of Operators

The main problem in sampling theory is to reconstruct a function from the values (samples) on some discrete subset of its domain. However, taking samples on an appropriate sampling set is not always practical or even possible|it can be that measuring devices are too expensive or scarce.

In the dynamical sampling it is assumed that the sparseness of the sampling locations can be compensated by involving dynamics. For example, when f is the initial state of a physical process (say, change of temperature or air pollution), we can iteratively sample its values at the same sampling locations as the time progresses, and try to recover f from the combination of these spatio-temporal samples.

We find conditions on the evolution operator and the sampling vectors under which the reconstruction of f is unique, stable etc. Some examples, as well as, recent results and open questions will be presented.

Mr. Viktor Reshniak: Reducing Computational Cost of the Multilevel Monte Carlo Method by Selection of Suitable Pathwise Integrators

The multilevel Monte Carlo (MLMC) method is a variance reduction technique for the efficient estimation of expected values of the quantities of interest associated with solutions of stochastic and random differential equations. By combining the ideas of multigrid discretization and Monte Carlo sampling, it allows to achieve the optimal asymptotical complexity of the estimator for the very large class of problems. The actual cost of the estimator, however, is more problem and solver dependent as the method requires one to solve a large number of decoupled deterministic problems. The efficiency of the estimator is hence strongly influenced by the complexity of the corresponding pathwise integrators. It is the task of this work to study several problems for which the significant reduction in the computational complexity of the MLMC estimator can be achieved by the appropriate problem and level dependent choice of deterministic solvers. Two particular problems are considered: integration of stiff SDEs and boundary values problems in randomly perturbed domains. For the considered problems, the optimal choice of pathwise integrators is discussed and both qualitative and quantitave analysis of the obtained computational improvements is performed.

Mr. Anton Dereventsov: Greedy Algorithms in Approximation Theory and Convex Optimizaiton

Greedy algorithms are designed to approximate an element f of a Hilbert (or Banach) space X by a linear combination of elements of a fixed set D (dictionary). An approximation produced by a greedy algorithm is naturally sparse with respect to a dictionary D, which makes greedy algorithms useful in many applications (e.g. compressed sensing or image processing). It was shown recently that greedy algorithms can be adapted to minimize a convex function E defined on a

Events (Continued)

Hilbert (or Banach) space X; the produced minimizer in this case is also sparse with respect to the selected dictionary D. In this talk we will discuss some types of greedy algorithms in approximation theory and convex optimization, their convergence results and practical implementation of these algorithms.

Dr. Vitalii Starchenko: Simulation of Fracture Dissolution in Rocks using OpenFOAM Toolkit

Typically, numerical models of fracture dissolution are simplified and represented by the system of one or two-dimensional fields. These approaches can explain many observations in nature, like the estimation of the breakthrough time or formation of the wormholes (highly localized flow paths). However, such modeling is limited to the early stages of fracture dissolution where the variations in aperture are small. As the dissolution patterns develop, a 2D model is unable to correctly capture the flow in the vicinity of the wormholes and tends to overestimate the degree of competition between adjacent flow paths.

This seminar will be devoted to the 3D modeling of fracture dissolution using the OpenFOAM toolkit. I will give an overview of the approach we used for the simulation and discuss the different limitations of two and three dimensional models. I will show how we used OpenFOAM's machinery to implement mesh relaxation, which is necessary for the stability of the solver, and boundary conditions for non-linear surface kinetics. Finally, I will briefly discuss future directions for the dissolution simulations in three dimensions.

Mr. Xuping Xie: Large Eddy Simulation Reduced Order Models

A large eddy simulation reduced order model (LES-ROM) framework for the numerical simulation of convection-dominated flows will be introduced. In this LES-ROM framework, the proper orthogonal decomposition (POD)is used to define the ROM basis and a ROM differential filter is used to define the large ROM structures. An approximate deconvolution (AD) approach is used to solve the ROM closure problem and develop a new AD-ROM. This AD-ROM is tested in the numerical simulation of the three-dimensional flow past a circular cylinder at a Reynolds number Re = 1000. Also, another novel data driven ROM method has been derived under the LES-ROM framework.

CSMD Seminars

Wayne Joubert: ORNL MiniApps Webinar Series: MiniSweep - A Proxy Application for Sn Radiation Transport Calculations

Abstract: Minisweep is a miniapp that models the performance of the sweep operation of the Denovo radiation transport code, part of the Exnihilo code suite used by CASL and current INCITE application. It is used for performance evaluation and porting of the Sn transport sweep algorithm to new computer architectures, currently supporting OpenMP, CUDA for GPUs and compiler directives for Intel Xeon Phi. In this talk we describe the structure of the underlying algorithm, techniques used to map the algorithm to advanced architectures, build and execution options and peformance characteristics of the code.

Dr. Swaroop Pophale: Designing next-generation programming models using lessons learned from PGAS

Abstract: With the coming of the exascale era and rapidly improving hardware, applications need well designed programming models that are simple to use but that capture key performance characteristics of next-generation systems. A programming model with the right abstractions also helps advance performance portability. Many lessons can be learned from Partitioned Global Address Space (PGAS) models that have explicit concepts to express data locality and control the affinity of the computations. As we move forward more PGAS concepts are making its way to mainstream programming models like MPI, OpenMP / OpenACC, etc. and the key challenges are to find right abstractions without increasing their complexity. This talk focuses on lessons learnt from PGAS and the new features being incorporated into mainstream models (OpenMP and OpenSHMEM) to tackle the problems of data locality, thread affinity and massive amounts of thread-level parallelism to improve the utilizations of memory, heterogeneous cores, and network interconnects.

Dr. Jean-Luc Fattebert: Scientific computing adventures in materials science

Abstract: Computer simulations are widely used by physicists, chemists, biologists and materials scientists to better understand the behavior and properties of matter, from the atomistic to the macroscopic levels. Dr. Fattebet will talk about a few numerical methods that he worked on over the years to speed up these simulations and improve their fidelity. From First-Principles Molecular Dynamics to Phase-Field models, Dr. Fattebert will describe how some advanced numerical algorithms can help application scientists efficiently use our most powerful High Performance Computing resources. While some challenges are inherent to the physical models and their specific equations, the different point of views, languages (in software and papers!), and jargons used by various

scientific communities, they are not to be underestimated when it comes to make good use of modern applied mathematics in real applications. It is often what makes the difference between a successful project with useful outputs, and an academic exercise with little use in real applications.

Chongxiao "Shawn" Cao: Utilization and Extension of Task-based Runtime for High Performance Dense Linear Algebra Applications

Abstract: On the road to Exascale computing, dynamic task-based runtimes can alleviate the disparity between hardware peak performance and application performance, by providing executions that unfold only based on the dataflow between tasks. In this presentation, I would like to introduce two parts of my PhD work related to task-based runtime. The first part is the design of a unified framework to run high-performance dense linear algebra applications for platforms equipped with multi-GPUs and multi-Xeon Phi coprocessors. A lightweight task-based runtime is utilized to manage the resource-specific workload, and to control the dataflow and parallel execution in hybrid system. The second part of this presentation is to introduce the fault tolerant design for a task-based runtime. Three additions to a dynamic task-based runtime have been explored to build a generic framework providing soft error resilience, including sub-DAG method, data logging method and algorithm-based fault tolerant method. We also take one step further to improve the general data logging method to a remote version to provide resilience for hard error.

Mark Berrill: ORNL MiniApps BlueJeans Seminar Series: The XRayTrace MiniApp

Abstract: This talk will cover the XRayTrace miniapp. XRayTrace is a program designed to simulate x-ray lasers that occur in laser-created plasmas. It solves the couples ray equations, intensity equations, and atomic physics to accurately predict the resulting x-ray laser properties including energy, near-field and far-field beam patterns, and the frequency and temporal shapes. The miniapp is designed to test the ray-trace behavior of the application under different architectures and programming models. The talk will cover a brief introduction into the design of X-ray lasers, the XRayTrace program, and finally the miniapp design. The design goals and implementation of the miniapp as well as a review of the existing parallel programming models implemented will be covered.

Bruno Turcksin: Matrix-free operator evaluation for additive manufacturing simulation

Abstract: There has been a long trend, in computer architecture, where the number of flops increases much faster than the amount and the bandwidth of memory. Today, even storing the matrix associated with the system can require a lot of the available memory. In many modern simulations, the use of matrices is only through sparse matrix-vector multiplications (SpMV) within Krylov solvers. However, SpMV is heavily bandwidth limited, both on CPU and GPU systems. GPUs have a much higher memory bandwidth than CPUs, and GPUs also have greater arithmetic capabilities. Moreover, GPUs have much less memory per core than CPUs. These factors favor stopping the separation of linear algebra from finite elements assembly routines and, instead, computing the action of the operator on the vector on the fly.

In this talk, Dr. Turcksin will discuss the implementation of the matrix-free operator evaluation in the context of additive manufacturing (AM) simulation. In particular, he will focus on the fast and accurate simulation of the heat transfer and the phase changes between the powder, the liquid, and the solid states of a given material. Phase changes and temperature variations are fundamental to predict the mechanical properties of materials. Therefore, these simulations are a necessary first step to predict the behavior of objects built using AM.

Dr. Turcksin will also briefly talk about adaptive mesh refinement strategies for AM simulation. Because the heat sources used in AM are very localized, it is important to use a mesh tailored to the problem solved. However, due to the time dependent nature of the problem, it is not possible to construct a single mesh that can be used throughout the whole simulation and therefore, the mesh needs to be adapted automatically.

George J. Nelson: Multiscale Transport in Energy Conversion and Storage Devices

Abstract: Contemporary energy storage and conversion devices are inherently multiscale, functional material systems. The operation of electrochemical energy storage and conversion devices relies upon composite material systems that provide both sites for chemical reactions and pathways for charge and reactant transport. Forging a stronger understanding of the multiscale interaction between transport phenomena, device and mesoscale geometry, and device performance and reliability can further advance the development of next generation storage and conversion technologies. To this end, x-ray and neutron imaging methods provide unique capabilities for direct observation of multiscale geometry, physics-based simulation using real device structures, and direct observation of material changes during operation. Three cases are addressed to illustrate these capabilities: x-ray tomography of Li-ion battery electrodes, mesoscale modeling of battery materials based on tomographic data, and in situ neutron imaging of enzymatic batteries. First, 3D x-ray nanotomography (XNT) and microtomography (µCT) of Li-ion battery electrode materials are presented for high capacity alloy anode materials and transition metal oxide cathode materials. Methods for elemental and chemical mapping within these systems are demonstrated. Absorption contrast imaging is presented as a means of

Events (Continued)

structural and elemental mapping within battery anodes and cathodes. X-ray absorption near edge structure (XANES) nanotomography is then presented as a means of tracking chemical variations related to charge and discharge. Following demonstration of current XNT and µCT work, mesoscale modeling of Li-intercalation is presented using Li-ion cathode XNT data as a computational domain. Observations from this model are complemented by microstructural analysis that permits description of cathode charge and discharge behavior in terms of dimensionless parameters. Finally, in situ and in operando imaging efforts are discussed in the context of neutron imaging performed on enzymatic batteries. The relevance of these methods to other electrochemical energy conversion and storage devices is discussed.

Eric Lingerfelt: BEAM: An HPC Pipeline for Nanoscale Materials Analysis and Neutron Data Modeling

Abstract: The Bellerophon Environment for Analysis of Materials (BEAM) enables scientists at ORNL's Center for Nanophase Materials Sciences and Spallation Neutron Source to leverage the integrated computational and analytical power of ORNL's Compute And Data Environment for Science (CADES) and the Oak Ridge Leadership Computing Facility to perform near real-time scalable analysis and modeling. At the core of this computational workflow system is a web and data server located at CADES that enables multiple, concurrent users to securely upload and manage data, execute materials science analysis and modeling workflows, and interactively explore results through custom visualization services. BEAM's long-term data management capabilities utilize CADES' petabyte-scale file system and enable users to easily manipulate remote directories and uploaded data in their private data storage area as if they were browsing on a local workstation. In addition, the framework facilitates user workflow needs by enabling integration of advanced data analysis algorithms and authenticated, "push-button" execution of dynamically generated workflows employing these algorithms on Titan, Eos, and Rhea at OLCF, as well as compute clusters at CADES. We will demonstrate band excitation analysis, principal component analysis, and de-noising of SPM and STEM data using a variety of HPC implementations including FORTRAN, R with pdbR, GPU-accelerated C++, and Java with Apache Spark – all tightly bound with parallel HDF5. In addition, we will discuss initial implementations of near real-time optimization and modeling of inelastic and quasi-elastic neutron scattering data utilizing Titan, CADES, and BEAM.

Community Service

- Rizwan Ashraf, program committee, ACM International Conference on Computing Frontiers 2017 (CF17)
- David E. Bernholdt, peer reviewer, DOE SBIR program
- Christian Engelmann, program committee, IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2017)
- Christian Engelmann, program committee, 4th Mathematical Methods and Algorithms for Extreme Scale Workshop (MATH-EX 2017)
- Christian Engelmann, peer reviewer, DOE Early Career program
- Saurabh Hukerikar, peer reviewer, Future Generation Computer Systems (FGCS)
- Jungwon Kim, Reviewer, Journal of Parallel and Distributed Computing
- Seyong Lee, Program Committee, IEEE International Parallel and Distributed Processing Symposium (IPDPS 2017)
- Seyong Lee, Program Committee, Posters and Research Demos track, IEEE/ACM International Symposium on Cluster, Cloud and Grid (CCGrid 2017)
- Seyong Lee, Program Committee, International Workshop on Accelerators and Hybrid Excascale Systems (AsHES 2017)
- Seyong Lee, Reviewer, TWMS Journal of Pure and Applied Mathematics
- Seyong Lee, Reviewer, Journal of Parallel Computing (ParCo)
- Seyong Lee, Reviewer, IEEE Micro
- Seyong Lee, Reviewer, Journal of Earth Science
- Seyong Lee, Reviewer, International Journal of High Performance Computing and Networking (IJHPCN)
- Tonglin Li, peer reviewer, IEEE Transactions on Computers
- Sarah Powers, panelist, Nebraska Conference for Undergraduate Women in Mathematics (NCUWM 2017)
- Philip C. Roth and R. Shane Cannon, guest editors, Parallel Computing: Systems & Applications 61(2017), Special Issue on Data-Intensive Scalable Computing Systems
- Sarat Sreepathi, Reviewer, IEEE Transactions on Parallel and Distributed Systems
- Geoffroy Vallee, peer reviewer, ORNL SEED
- Geoffroy Vallee, program committee, 7th International Conference on Cloud Computing and Services Science (CLOSER 2017)
- Jeffrey S. Vetter, Program Committee, IEEE International Parallel and Distributed Processing Symposium (IPDPS 2017)

Community Service (Continued)

- Jeffrey S. Vetter, Chair, Research Posters, International Supercomputer Conference (ISC).
- Jeffrey S. Vetter, Program Committee, International Symposium on Performance Analysis of Systems and Software (ISPASS)
- Jeffrey S. Vetter, Program Committee, Symposium on Principles and Practice of Parallel Programming (PPoPP)
- Jeffrey S. Vetter, Co-Guest Editor, IEEE Micro Special Issue on Post Moores Era Supercomputing.
- Gregory Watson, program committee, Tools for Program Development and Analysis in Computational Science (TOOLS 2017)

Awards and Recognition

Peter Jantsch - NSF Fellowship

Peter Jantsch has been awarded the prestigious National Science Foundation's Mathematical Sciences Postdoctoral Research Fellowship (MSPRF).

The purpose of the MSPRF is to support future leaders in mathematics and statistics by facilitating their participation in postdoctoral research environments that will have maximal impact on their future scientific development. There are two options for awardees: Research Fellowship and Research Instructorship. Awards support research in areas of mathematics and statistics, including applications to other disciplines.

Peter's advisior is Computational and Applied Mathematics Group leader - Clayton Webster.

Publications of Note

Durango: Scalable Synthetic Workload Generation for Extreme-Scale Application Performance Modeling & Simulation

Christopher D. Carothers, Jeremy S. Meredith, Mark P. Blanco, Jeffery Vetter, Misbah Mubarak, Justin LaPre, and Shirley Moore ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS), Singapore, May 2017. DOI: http://dx.doi.org/10.1145/3064911.3064923

Performance modeling of extreme-scale applications on accurate representations of potential architectures is critical for designing next generation supercomputing systems because it is impractical to construct prototype systems at scale with new network hardware in order to explore designs and policies. However, these simulations often rely on static application traces that can be difficult to work with because of their size and lack of flexibility to extend or scale up without rerunning the original application. To address this problem, this paper presents a new technique for generating scalable, flexible workloads from real applications. A prototype called Durango combines a proven analytical performance modeling language, Aspen, with the massively parallel HPC network modeling capabilities of the CODES framework. The models are compact, parameterized and representative of real applications with computation events. They are not resource intensive to create and are portable across simulator environments. This papers demonstrate the utility of Durango by simulating the LULESH application in the CODES simulation environment on several topologies and shows that Durango is practical to use for simulation without loss of fidelity, as quantified by simulation metrics. Additionally, the paper demonstrates the efficacy of Durango's direct integration approach, which links Aspen into CODES as part of the running network simulation model. Here, Aspen generates the application-level computation timing events, which in turn drive the start of a network communication phase. Results show that Durango's performance scales well when executing both torus and dragonfly network models on up to 4K Blue Gene/Q nodes using 32K MPI ranks, Durango also avoids the overheads and complexities associated with extreme-scale trace files.

ASA-FTL: An adaptive separation aware flash translation layer for solid state drives

W. Xie, Y. Chen, P.C. Roth; Parallel Computing 61:3-17, (2017); http://dx.doi.org/10.1016/j.parco.2016.10.006

The flash-memory based Solid State Drive (SSD) presents a promising storage solution for increasingly critical data-intensive applications due to its low latency (high throughput), high bandwidth, and low power consumption. Within an SSD, its Flash Translation Layer (FTL) is responsible for exposing the SSD's flash memory storage to the computer system as a simple block device. The FTL design is one of the dominant factors determining an SSD's lifespan and performance. To reduce the garbage collection overhead and deliver better performance, we propose a new, low-cost, adaptive separation-aware flash translation layer (ASA-FTL) that combines sampling, data clustering and selective caching of recency information to accurately identify and separate hot/cold data while incurring minimal overhead. We use sampling for light-weight identification of separation criteria, and our dedicated selective caching mechanism is designed to save the limited RAM resource in contemporary SSDs. Using simulations of ASA-FTL with both real-world and synthetic workloads, we have shown that our proposed approach reduces the garbage collection overhead by up to 28% and the overall response time by 15% compared to one of the most advanced existing FTLs. We find that the data clustering using a small sample size provides significant performance benefit while only incurring a very small computation and memory cost. In addition, our evaluation shows that ASA-FTL is able to adapt to the changes in the access pattern of workloads, which is a major advantage comparing to existing fixed data separation methods.

About CSMD

The Computer Science and Mathematics Division (CSMD) is ORNL's premier source of basic and applied research in high-performance computing, applied mathematics, and intelligent systems. Basic and applied research programs are focused on computational sciences, intelligent systems, and information technologies.

Our mission includes working on important national priorities with advanced computing systems, working cooperatively with U.S. industry to enable efficient, cost-competitive design, and working with universities to enhance science education and scientific awareness. Our researchers are finding new ways to solve problems beyond the reach of most computers and are putting powerful software tools into the hands of students, teachers, government researchers, and industrial scientists.

The Division is composed of seven Groups. These Groups and their Group Leaders are:

- Computational Applied Mathematics – Clayton Webster
- Computational Data Analytics – Tom Potok
- Computer Science Research – David Bernholdt
- Data System Sciences & Engineering – Tommy Nelson
- Discrete Computing Systems – Kalyan Perumalla
- Future Technologies – Jeff Vetter
- Scientific Data – Scott Klasky

About this Newsletter

This newsletter is compiled from information submitted by CSMD Group leaders, public announcements and searches.
Please contact Daniel Pack if you have information you would like to contribute.

Contact Information

CONTACTS

CSMD Director - Barney Maccabe - maccabeab@ornl.gov
Division Secretary - Lora Wolfe - wolfelm@ornl.gov
Division Finance Officer - Ursula Henderson - hendersonuf@ornl.gov
Technical Communications - Daniel Pack - packdl@ornl.gov

LINKS

Computer Science and Mathematics Division - www.csm.ornl.gov
Computing and Computational Sciences Directorate - computing.ornl.gov
Oak Ridge National Laboratory - www.ornl.gov
Computational and Applied Mathematics Group - www.cam.ornl.gov
Future Technologies Group - www.ft.ornl.gov

Newsletter Archive - www.csm.ornl.gov/newsite/newsletters.html

OAK RIDGE National Laboratory