ISSUE
23
WINTER
2016

QUARTERLY NEWSLETTER FOR THE COMPUTER SCIENCE AND MATHEMATICS DIVISION

CSMD

this issue

Science Highlights P.01
News P.14
Community Service P.14
Events P.14
Awards and Recognition P.15
New Faces P.17
About CSMD P.20

Science Highlights

Selective Electrochemical Conversion of Carbon Dioxide to Ethanol

Yang Song, Rui Peng, Dale Hensley, Peter V. Bonnesen, Liangbo Liang, Zili Wu, Harry M. Meyer III, Miaofang Chi, Bobby Sumpter, and Adam J. Rondinone

Researchers were able to use a nanostructured catalyst to electrochemically convert dissolved carbon dioxide to ethanol under ambient conditions with 63% yield.

This catalyst achieves unprecedented selectivity for a complex reaction and demonstrates that nanoscale texture can be used to control reaction outcomes.

This work highlights an electrochemical process based on tiny spikes of a carbon/nitrogen nanocomposite and copper nanoparticles to turn electricity and carbon dioxide, a greenhouse gas, into ethanol. By arranging common materials that avoid traditional catalysis such as expensive or rare metals, into highly nanotextured structure, unwanted chemical side reactions are limited, providing a high yield (63%) of the desired product (84%), ethanol. Electrochemical analysis and density functional theory calculations suggest that a preliminary mechanism driving the direct reduction process involves active sites on the Cu nanoparticles and the carbon nanospikes working in tandem to control the electrochemical reduction of the carbon monoxide dimer to alcohol.

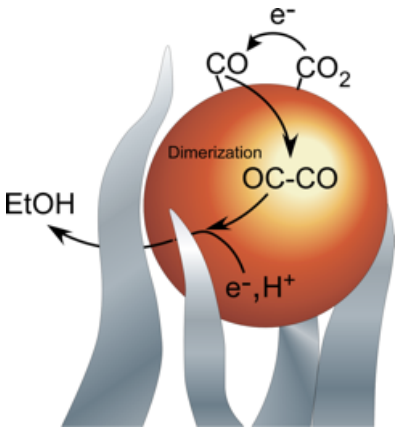

Copper nanoparticles imbedded within nitrogen-doped carbon nanospikes catalyze the electroreduction of CO2 directly to ethanol. The final electrochemical reduction step is believed to take place at the nanoparticle/nanospike interface.

This research was conducted at the Center for Nanophase Materials Sciences, which is a Department of Energy (DOE) Office of Science User Facility.

Publication: Yang Song, Rui Peng, Dale Hensley, Peter V. Bonnesen, Liangbo Liang, Zili Wu, Harry M. Meyer III, Miaofang Chi, Bobby Sumpter, Adam J. Rondinone, Direct electrochemical reduction of CO2 to Ethanol using a multiple-site, nanostructured Cu/Carbon electrode, ChemSelect 1, 1-8 (2016). DOI: 10.1002/slct.201601169

OAK RIDGE National Laboratory

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.

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.

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

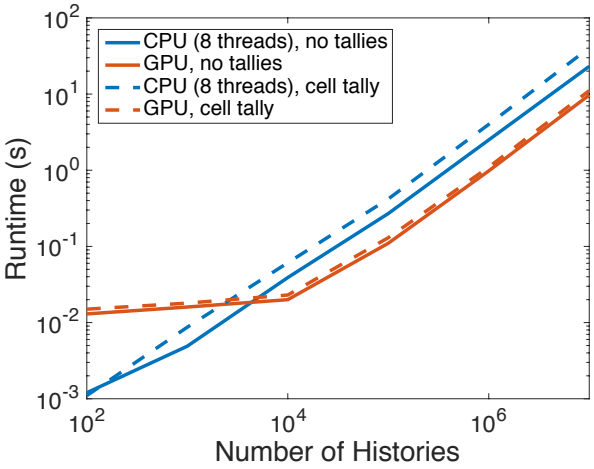

Comparison of CPU and GPU runtimes with and without tallies enabled for an increasing number of neutrons in the simulation. Each thread computes the trajectory of a neutron from birth to death as it interacts with the surrounding material in a nuclear reactor. Tallies allow for the measurement of neutron 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.

Science Highlights

Probing Multiscale Transport and Inhomogeneity in a Lithium-Ion Pouch Cell Using In-Situ Neutron Methods

H. Zhou, K. An, S. Allu, S. Pannala, J. Li, H. Z. Bilheux, S. K. Martha, J. Nanda

A powerful nondestructive tool that can be utilized to understand the multiscale nature of lithium transport is developed with combined in-situ radiography and diffraction supported by 3D multiscale electrochemical modeling.

This work improves our ability to effectively analyze in-homogeneities of a thick electrodes during various states of lithiation at a cell level that are important for increasing energy density of batteries.

We demonstrate the lithiation process in graphitic anodes using in-situ neutron radiography and diffraction in a single-layer pouch cell. The variation in neutron absorption contrast in graphite shows a direct correlation between the degree of lithiation and the discharge potential. The experimental neutron attenuation line profiles across the graphite electrode at various discharge times (potentials) were compared with lithium concentration profiles computed using a 3D electrochemical transport model. In conjunction with imaging/radiography, in-situ neutron diffraction was carried out to obtain information about the local structural changes during various stages of lithiation in carbon. Combined in-situ radiography and diffraction supported by 3D multiscale electrochemical modeling opens up a powerful nondestructive tool that can be utilized to understand the multiscale nature of lithium transport as well as observe various in-homogeneities at a cell level.

This research used resources at the High Flux Isotope Reactor and Spallation Neutron Source, a DOE Offi ce of Science User Facility operated by the Oak Ridge National Laboratory. The authors acknowledge support from the Office of Vehicle Technology, EERE, DOE.

Publication: H. Zhou, K. An, S. Allu, S. Pannala, J. Li, H. Z. Bilheux, S. K. Martha, J. Nanda, "Probing Multiscale Transport and Inhomogeneity in a Lithium-Ion Pouch Cell Using In-Situ Neutron Methods" ACS. Energy Letters 1(5), 2016; http://dx.doi.org/10.1021/acsenergylett.6b00353

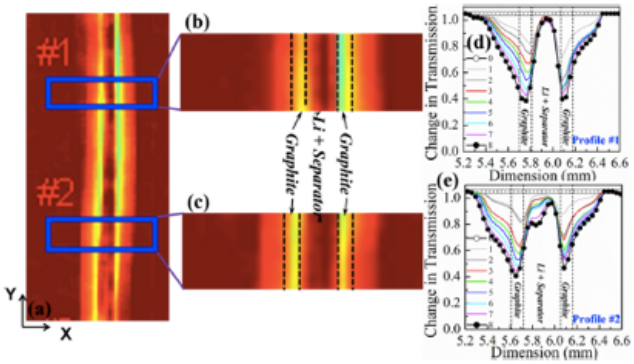

Line profile analyses at two different cell heights in the Y direction: (a) position of the two line profiles used for transmission as a function of the X direction studies; (b) neutron image of position #1 at the discharge state of point 4; (c) neutron image of position #2 at the discharge state of point 4; (d) line profile #1 transmission plot; and (e) line profile #2 transmission plot.

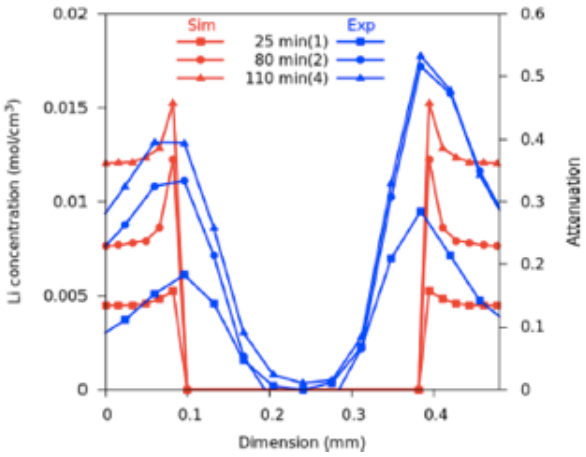

Comparison of the calculated lithium concentration as a function of electrode thickness (red curves) to the experimental line profiles of the attenuation (blue curves) at discharge times of 25, 80, and 110 min

Science Highlights (Continued)

Resilience Design Patterns - A Structured Approach to Resilience at Extreme Scale

S. Hukerikar, and C. Engelmann

Researchers developed a structured approach for managing high-performance (HPC) resilience using design patterns. Developed a complete specification of resilience design patterns, including a pattern catalog, classification scheme and a design framework to build resilience solutions using the patterns.

The resilience design patterns specification provides HPC system architects and application developers with a collection of reusable design elements, and a discipline to combine an essential set of design patterns into productive and efficient resilience solutions. The structured approach enables the systematic improvement of resilience in HPC systems, keeping scientific applications running to a correct solution in a timely and efficient manner in spite of frequent faults, errors, and failures.

A design pattern describes a generalizable solution to a recurring problem that occurs within a well-defined context. Patterns are often derived from best practices used by designers and they contain essential elements of the problems and their solutions. They provide designers with a template on how to solve a resilience problem that may be used in many different situations. The patterns in the catalog provide solutions for detection, containment and recovery for faults, errors and failures. This work also defined a conceptual framework based on the notion of design spaces that guides hardware and software designers and architects, as well as application developers, in navigating the complexities of developing effective resilience solutions. The framework enables designers to combine the patterns and refine their interactions to create alternative resilience solutions for a specific problem, each with different efficiency and complexity characteristics.

This work was performed at Oak Ridge National Laboratory (ORNL). It was sponsored by the US Department of Energy Office of Science Early Career Research Program.

Publications:

• Saurabh Hukerikar and Christian Engelmann. Resilience Design Patterns: A Structured Approach to Resilience at Extreme Scale (Version 1.0). Technical Report, ORNL/TM-2016/687, Oak Ridge National Laboratory, Oak Ridge, TN, USA, October, 2016.

• Saurabh Hukerikar and Christian Engelmann. Resilience Design Patterns: A Structured Approach to Resilience at Extreme Scale (Version 1.0). Technical Report, ORNL/TM-2016/687, Oak Ridge National Laboratory, Oak Ridge, TN, USA, October, 2016.

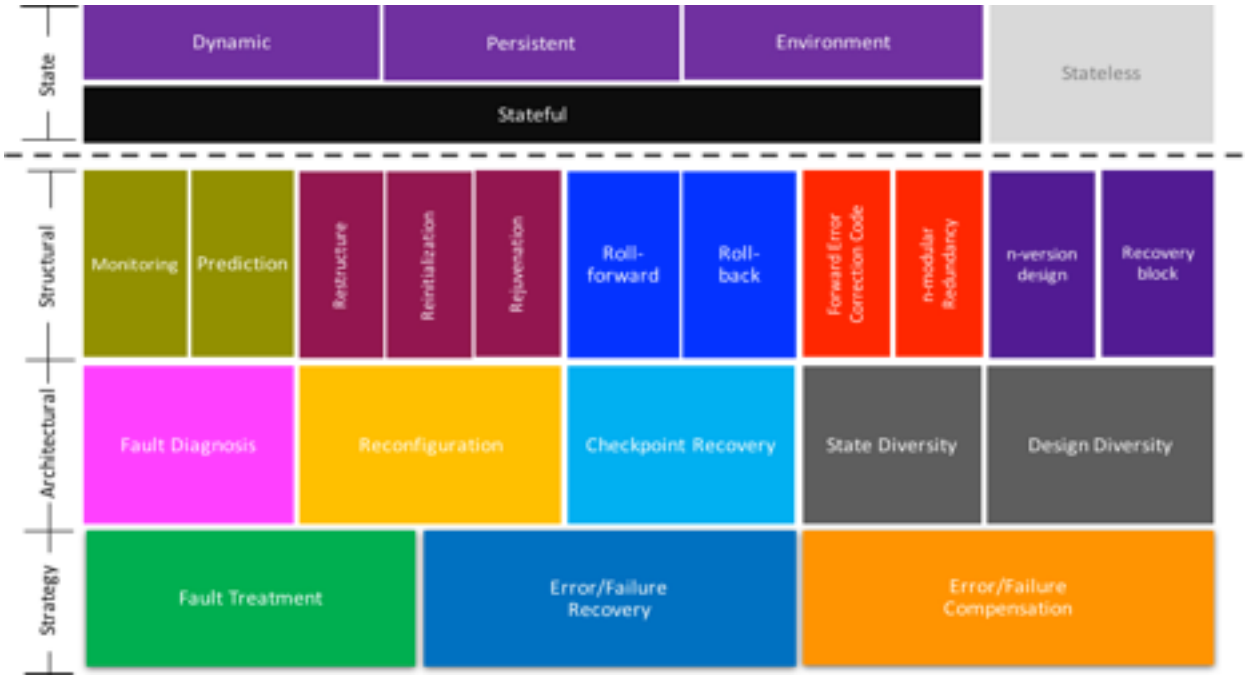

Classification of Resilience Design Patterns

Science Highlights (Continued)

Improving Computational Efficiency of Xolotl for Modeling Gas Bubble Evolution in Nuclear Materials

Philip C. Roth (ORNL), David Bernholdt (ORNL), Sophie Blondel (UT-K), Brian Wirth (UT-K)

Optimization of Xolotl reaction-diffusion cluster dynamics code reduced reaction network initialization time for nuclear fuel applications from 23+ hours to under 7 seconds (improvement of 104 times)

This work enables nuclear fuel simulations that were previously intractable due to extreme reaction network initialization times. The work benefits Xolotl users conducting both nuclear fusion simulations and plasma surface interaction simulations.

Xolotl is a reaction-diffusion cluster dynamics code developed within the Plasma Surface Interactions (PSI) SciDAC project and adapted for nuclear fuel simulation in the NE-SciDAC-Pilot project with collaboration from the SciDAC Institute for Sustained Performance, Energy, and Resilience (SUPER). Prior to this work, Xolotl's reaction network initialization did not scale well with respect to problem size. Its performance was adequate for smaller PSI reaction networks, but the desired NE test problem scales were intractable due to poor reaction network initialization scalability. As part of the PSI-SUPER collaboration, a lightweight performance data collection infrastructure was designed and implemented within Xolotl using ASCR funding. In this work, we used this performance infrastructure to diagnose performance problems with Xolotl's reaction network initialization. We identified the primary cause of this performance problem to be the use of inappropriate data structures (maps of reactant compositions) as keys in reaction network lookup maps. We modified Xolotl to use strings with a canonical format as reaction network lookup map keys because they are much cheaper to compare than the previous reactant composition maps. Our use of string keys also enabled reuse for subsequent optimizations. Our second major Xolotl optimization addressed the performance problem of inefficient removal of reactants from the reaction network during the formation of supercluster reactants. Supercluster reactants are reactants that are used to represent multiple individual reactants so as to control the simulation's memory footprint and computational cost while retaining the simulation's fidelity. Using Xolotl's built-in performance infrastructure, we identified redundant examination of the collection of known reactants during reactant removal as a significant performance problem. By modifying the code to remove reactants in bulk as opposed to one-at-a-time, by ensuring that each item in the known reactant collection was examined exactly once when deciding to remove it, and by leveraging our use of string keys to represent reactant composition, we achieved an additional significant performance improvement compared to the initial Xoltol implementation.

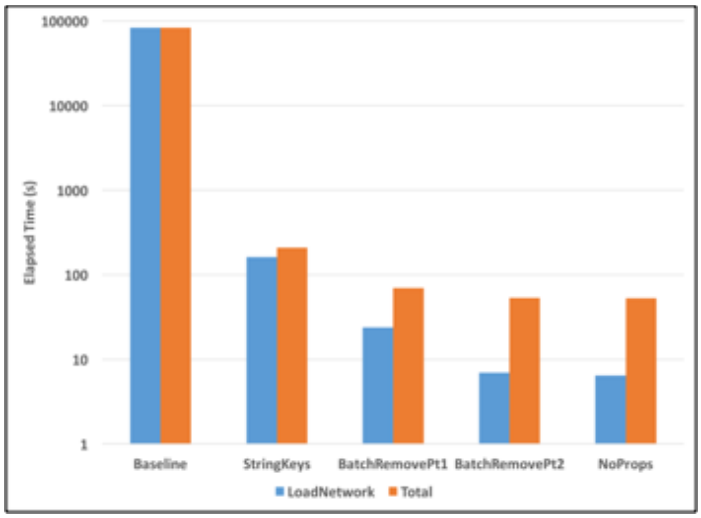

Reaction network initialization time (blue) and total run time (orange) for NE test problem, performed on OLCF's Eos system. Left-to-right bars indicate performance of initial version to optimized versions. Note elapsed times are plotted on log scale.

This work is funded by DOE office of Advanced Scientific Computing Research (ASCR) and used resources of the Oak Ridge Leadership Computing Facility (OLCF)

Science Highlights (Continued)

Application of Data Analytics to Additive Manufacturing

S. Powers, R. Dehoff, V. Paquit, C. Steed and D. Kistler

Researchers developed both static and streaming data heuristics for analytically making recommendations about when to halt an advanced manufacturing build as well as identifying build areas for further investigation.

Current trends in advanced manufacturing largely rely on microstructural analysis of builds in order to improve the manufacturing process. This work demonstrates that a multipronged approach which includes the use of data analytics can provide additional insight into failures. The research provides tools to automate precursor steps and hone in on areas of concern, thereby reducing the analysis time and pinpointing areas of concern.

Recent advances in additive manufacturing have led to many success stories of large 3D printed objects (e.g., the Shelby Cobra car) and leave the industry poised for rapid growth. One of the many challenges is the certification process, which currently must be done for each part. During the build progression, many parameters have the potential to influence the formation of defects as well as the final geometry. While a large focus has been placed on understanding the microstructural properties from a material science perspective, the vast quantities of data generated during the build have largely been unexploited. This work describes a multi-pronged approach for data discovery, engaging multiple analytic tools as well as a framework to ingest and house the data itself in an effort to identify areas for potential improvement and promote the potential for advanced defect detection.

This work was sponsored by the Department of Energy and was performed at the Oak Ridge National Laboratory.

Publication: S. Powers, R. Dehoff, V. Paquit, C. Steed and D. Kistler, "Application of Data Analytics to Additive Manufacturing," in Proc. of 2016 INFORMS workshop on Data Mining and Decision Analytics (DM-DA 2016), Nashville, TN, USA, November 2016.

A DOE Research Laboratories View on Time Synchronization Needs and Challenges

T. Jones, J. Gracia, and P. Ewing

Researchers characterized and quantified barriers and gaps related to time synchronization and its application to the nation's power grid.

This work enhances our ability to effectively diagnose and operate the nation's power grid in times of normal operation as well as during natural or man-made induced interruptions.

Sensors distributed physically throughout the nation are an important aspect to the monitoring and operation of the power grid. These sensors, including the common Phasor Measurement Units (PMUs), do high-speed grid monitoring with time stamps dependent on time-synchronized to UTC with microsecond accuracy. Timing errors from the time source can cause incorrect synchrophasor data. Such errors can create false analytical conclusions and in the future could drive undesirable and possibly dangerous automated grid operations with synchrophasor-based controls. The power sector seeks to protect future grid operations with better timing tools and practices to improve robustness and resilience. PNT should be assumed unreliable at both source and receiving points, and new measures to assure accurate, reliable time stamps against multiple failure modes are needed.

This work was performed ORNL / sponsored by DOE EERE.

Through research and outreach, ORNL is helping the power grid community to implement consistent, reliable time across a substation and synchronized across a power system. Power system owners and operators use primarily GPS as the source for navigation, position and timing. If we lose GPS today, it will complicate (higher cost, longer duration, lower efficiency), but not kill, grid operations. However, GPS and alternate time sources (and the ways we deliver and use them) need to become more reliable for time-synchronized applications to become mission-critical in the future. ORNL's efforts seek to help our user community implement consistent, reliable time across a substation and synchronized across a power system.

The Power Grid (top) is becoming more complex with new requirements for bi-directional flow stemming from power generation at the residential and neighborhood levels. The new capabilities are bringing about increased requirements for dynamic sensor equipment.

Publication: Terry Jones, "A DOE Research Laboratories View on Time Synchronization Needs and Challenges", IEEE/NIST Timing Challenges in the Smart Grid Workshop, Gaithersburg, MD. Oct, 2016.

Science Highlights (Continued)

Benchmark Generation and Simulation at Extreme Scale

Mahesh Lagadapati (NCSU), Frank Mueller (NCSU), and Christian Engelmann (ORNL)

Researchers developed a new method for investigating the performance of parallel applications at scale on future high-performance computing architectures and the impact of different architectural choices.

This novel solution combines simulations using models of future HPC systems and communication traces of current HPC applications, such that benchmarks are automatically generated from communication traces and executed within a simulation. This technique provides a number of benefits, such as eliminating the data intensive trace replay, enabling simulations at different scales and obfuscating proprietary application implementation details.

Investigating the performance of parallel applications at scale on future architectures and the performance impact of different architectural choices is an important component of HPC hardware/software co-design. Simulations using models of future HPC systems and communication traces from applications running on existing HPC systems can offer an insight into the performance of future architectures. Instead of simply replaying a trace within a simulator, this work promotes the generation of a benchmark from traces. This benchmark is subsequently exposed to simulation using models to reflect the performance characteristics of future-generation HPC systems. This technique provides a number of benefits, such as eliminating the data intensive trace replay and enabling simulations at different scales. This work features novel software co-design aspects, combining the ScalaTrace tool to generate scalable trace files, the ScalaBenchGen tool to generate the benchmark, and the xSim tool to assess the benchmark characteristics within a simulator.

This work was performed at North Carolina State University (NCSU) and Oak Ridge National Laboratory (ORNL). It was sponsored by the US Department of Energy Office of Science, the US National Science Foundation, and ORNL's Laboratory Directed Research and Development program.

Publication: Mahesh Lagadapati, Frank Mueller, and Christian Engelmann. Benchmark Generation and Simulation at Extreme Scale. In Proceedings of the 20th IEEE/ACM International Symposium on Distributed Simulation and Real Time Applications (DS-RT) 2016, pages 9-18, London, UK, September 21-23, 2016. IEEE Computer Society, Los Alamitos, CA, USA. ISBN 978-1-5090-3506-9. ISSN 1550-6525. Acceptance rate 42.0% (21/50). Best paper candidate.

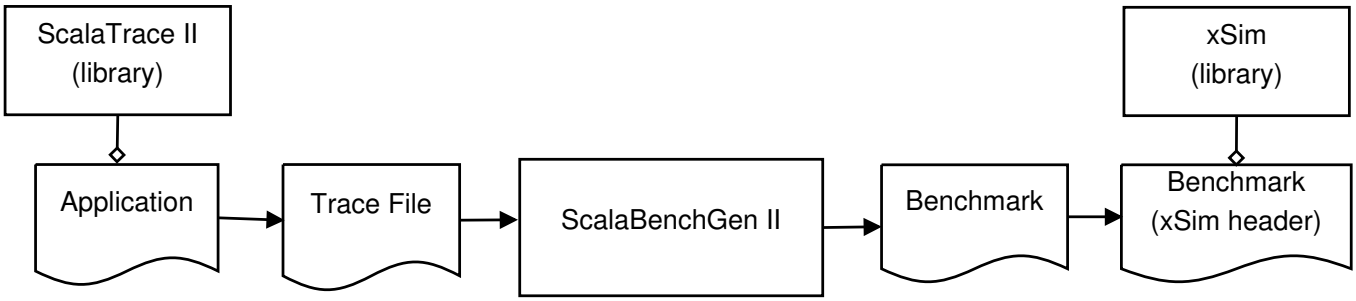

Benchmark Generation and Simulation

Science Highlights (Continued)

DISP: Optimized Initialization for Scalable MPI Startup

Prof. Suzanne Rivoire and her students (Sonoma State University), Chung-Hsing Hsu (ORNL)

Researchers developed a new scalable startup scheme with three internal techniques, namely Delayed Initialization, Module Sharing and Prediction-based Topology Setup (DISP). The DISP scheme greatly benefits the collective initialization of the Cheetah module and helps boost the performance of non-collective initialization in the Tuned module. Evaluation of our implementation on Titan with up to 4096 processes show that our delayed initialization can speed up the startup of Tuned and Cheetah by an average of 32.0% and 29.2%, respectively. The module sharing can reduce the memory consumption of Tuned and Cheetah by up to 24.1% and 83.5%, respectively, and our prediction-based topology setup can speed up the startup of Cheetah by up to 80%.

DISP is a suite of optimizations towards scalable MPI startup. In contrast to the conventional full-fledged, non-reusable, and communication-intensive initialization that has been revealed to be extremely costly in terms of both execution time and memory consumption, DISP initializes the communicators only partially, offloads the laborious part of initialization to where they are needed, reuse collective modules whenever possible, and only resort to collective communication for initialization when necessary. Our solution features a combination of three techniques. Firstly, we design a delayed initialization scheme that postpones the completion of communicator initialization to the point where the communicator is actually used. This replaces existing communicator creation routines where the communicators are fully initialized at the beginning; the delayed initialization only initializes a shadow communicator on top of a full- fledged communicator. The shadow communicator abstracts the essential subset of the communicator's information. The communicators are then initialized on an on-demand basis, only to be fully initialized when it is used for the first time by a collective operation. Secondly, we have both temporal and spatial module sharing mechanisms that realize the sharing of communicator modules between either different communicator objects of one process, or several processes that reside on the same node. The eligibility criteria of two communicators for module sharing are that their shadow communicators must be identical. Lastly, we improve the MPI startup with a mechanism called prediction-based topology setup that can locally compute most of the predictable hierarchical topology needed by Cheetah initialization without communicating back and forth over the network. Our design requires no changes to MPI APIs or the user application.

This work was performed at ORNL and Florida State University.

Publication: Huansong Fu, Swaroop Pophale, Manjunath Gorentla Venkata, and Weikuan Yu. 2016. DISP: Optimizations towards scalable MPI startup. In Proceedings of the First Workshop on Optimization of Communication in HPC (COM-HPC '16). IEEE Press, Piscataway, NJ, USA, 53-62. DOI: https://doi.org/10.1109/COM-HPC.2016.11

Science Highlights (Continued)

Tuning Defects in Monolayer Titanium Carbide (Ti3C2Tx) "Mxene" at the Atomic Level

X. Sang, Y. Xie, M.-W. Lin, M. Alhabeb, K. L. Van Aken, Y. Gogotsi, P. R.C. Kent, K. Xiao, and R. R. Unocic

Researchers showed, that by using Aberration-corrected STEM, that Ti vacancy density could be controlled by the etchant concentration during preparation of single-layer Ti3C2Tx MXene.

The presence of defects strongly influences materials properties; thus, a detailed understanding of the MXene single layer structure and point defects is foundational for unlocking physiochemical properties.

The 2D transition metal carbides or nitrides, Mxene phases, are emerging as a group of materials showing great promise in lithium ion batteries and supercapacitors. Until now, characterization and properties of single layer MXenes have been scarcely reported. Here, using scanning transmission electron microscopy (STEM), we determined the atomic structure of freestanding monolayer Ti3C2Tx flakes prepared via the minimally intensive layer delamination (MILD) method and characterized the different point defects that are prevalent in the monolayer flakes. We determine that the Ti vacancy concentration can be controlled by the etchant concentration during preparation. Density function theory-based calculations confirm the defect structures and predict that the defects can influence the surface morphology and termination

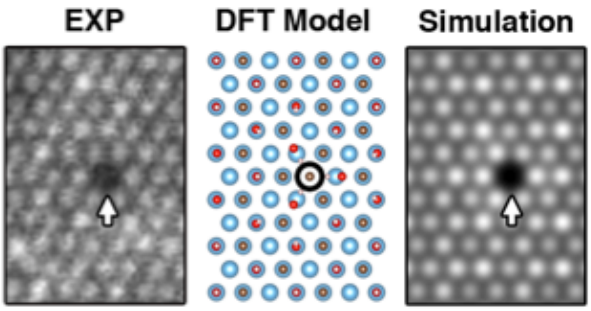

Comparison between experimental (EXP) HAADF-STEM image, defect crystal structure from DFT, and simulated HAADF-STEM image of a single Ti vacancy.

groups, but do not strongly influence the metallic conductivity. Using devices fabricated from single and few-layer Ti3C2Tx MXene flakes, the effect of the number of layers in the flake on conductivity has been demonstrated.

This work was conducted at the Center for Nanophase Materials Sciences, a U.S. Department of Energy (DOE) Office of Science user facility.

Publication: X. Sang, Y. Xie, M.-W. Lin, M. Alhabeb, K. L. Van Aken, Y. Gogotsi, P. R.C. Kent, K. Xiao, R. R. Unocic. ACS Nano, 10, 9193-9200 (2016). DOI: 10.1021/acsnano.6b05240

Phase Recognition from Power Traces of HPC Workloads

Prof. Suzanne Rivoire and her students (Sonoma State University), Chung-Hsing Hsu (ORNL)

Researchers developed a software method that is able to correctly detect and associate the phases of a power trace to known kernels with a high percentage.

This work addresses some of the limitations in the current state-of-the-art in understanding workloads' power consumption behavior. As power constraints and energy costs loom large in high performance computing, it can help schedulers to optimize performance under a power constraint.

Prior work has shown that power consumption traces of HPC workloads exhibit distinctive statistical characteristics, which allows the workload that generated a given power trace to be inferred with high accuracy. However, these power signatures apply to the entire power trace, with no ability to break it down further into phases or to recognize novel combinations of known workloads. In this work, we propose a software method for partitioning a power trace into phases and matching each phase to a known kernel or workload. We evaluated the method on a set of 388 power traces collected from 21 benchmarks, including CPU stressors, the NAS Parallel Benchmarks; and Mahout data analytics workloads. The results show that the method is able to, on average, attribute 78% of the points in a concatenated trace to the correct kernel.

This research was sponsored by the Department of Defense.

Publication: Joseph Granados, Jake Probst, Nick Armour, Jeff Bahns, Suzanne Rivoire, and Chung-Hsing Hsu, "Phase Recognition from Power Traces of HPC Workloads," 7th International Workshop in Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems, November 2016. (PTS#: 72322)

Science Highlights (Continued)

Stretching Semiconductors to Tune Optoelectronic Properties

J. Lee, J. Huang, B. G. Sumpter, and M. Yoon

Researchers use first-principles calculations to reveal how mechanical strain be used to enhance the optoelectronic properties of two-dimensional transition metal dichalcogenide (TMD) lateral heterostructures.

The work unambiguously shows that strain can simultaneously control the the bandgap, thus opening a route for tuning of optoelectronic properties in TMD lateral heterostructures.

In this work we report the influence of strain on the optoelectronic properties (power conversion efficiency) of TMD-based lateral heterostructures such as MoX2/WX2 (X=S, Se,Te), by first-principles density functional theory and quasiparticle GW calculations. The results unambiguously show that strain can simultaneously control the energetic distribution of bandgap, thus open a route for tuning of optoelectronic properties in TMD lateral heterostructures. Specifically we find that lateral MoX2/WX2 (X=S, Se, Te) heterostructures develop type II band alignment that facilitates efficient electron–hole separation for light detection/harvesting. In addition, the type II character is robust and can be preserved up to 12% of uniaxial strain. From the strain-dependent bandgap and the band offset, we find that uniaxial tensile strain can significantly increase the power conversion efficiency of lateral heterostructures. For example, 4% uniaxial tensile strain increases the power efficiency of both lateral MoSe2/WSe2 and MoS2/WS2 heterostructures by 15% up to 35%. We thus expect that these strained TMD lateral heterostructures can be utilized as novel architectures for efficient optoelectronic device development by selective control of the bandgap in energetic distribution.

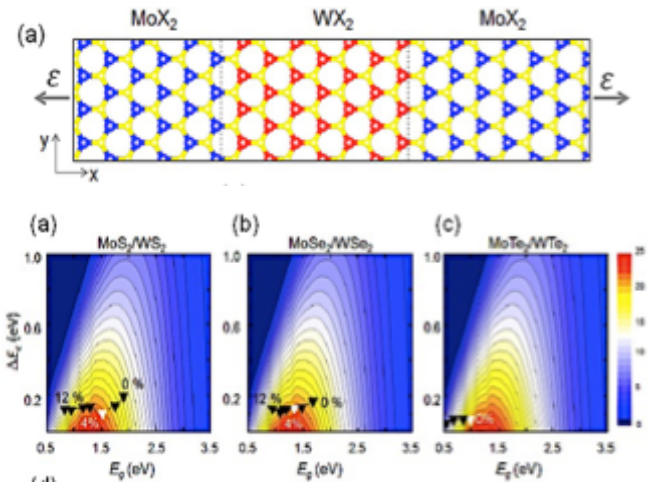

Top: (a) The lateral heterostructure, MoX2/WX2 (X=S, Se, Te), under a uniaxial strain ε.

Bottom: Computed power conversion efficiency contours as a function of the bandgap (Eg) of WX2 and conduction band offset (ΔEc) of WX2, where the computed Eg and ΔEc at each strain are overlaid on the contour.

This work was conducted at the Center for Nanophase Materials Sciences, a U.S. Department of Energy (DOE) Office of Science user facility. J.L. acknowledges support from ORNL Laboratory Directed Research and Development. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the US DOE (contract no. DE-AC02-05CH11231).

Publication: J. Lee, J. Huang, B. G. Sumpter, M. Yoon., 2D Materials (2016).

Science Highlights (Continued)

Engineering Materials for Ultrafast Charge Transfer

A. Boulesbaa, Kai Wang, Masoud Mahjouri-Samani, Mengkun Tian, Alexander A. Puretzky, Ilia Ivanov, Christopher M. Rouleau, Kai Xiao, Bobby G. Sumpter, and David B. Geohegan

Researchers show that following selective photo-excitation of the 2D-WS2, electron transfer to the 0D-QDs takes place on an ultrafast sub-45 fs time-scale.

The hybrid excitons are several times more weakly bound than in parent 2D materials, enabling dissociation that can improve the conversion efficiency of light into electricity.

In a synthesized hybrid nanostructure composed of a monolayer of tungsten disulfide (2D-WS2) and a single layer of core-shell cadmium selenide (CdSe)-zinc sulfide (ZnS) quantum dots (0D-QDs), we show that following selective photo-excitation of the 2D-WS2, electron transfer to the 0D-QDs takes place on an ultrafast sub-45 fs time-scale. The electron now located at the 0D-QDs is bound to the hole left at the 2D-WS2 forms a hybrid exciton that is several times more weakly bound than the tightly bound excitons in 2D-WS2 MLs. Such photoinduced interfacial charge transfer is at the heart of many applications, including photovoltaics, photocatalysis, and photodetection and thus may open new opportunities for optoelectronic applications of 2D and 0D nanomaterials.

Monolayer two-dimensional transition metal dichalcogenide (2D-TMD) semiconductors recently emerged as new candidates for atomically-thin optoelectronics. Unfortunately, excitons in these materials are tightly bound (>500 meV), which makes their dissociation challenging and constrains 2D-TMD application in photovoltaic and photodetection applications. In this work, a novel approach was advanced to mediate this problem by demonstrating how strongly-bound excitons of 2D tungsten disulfide (WS2) monolayers, covered with a single layer of semiconductor quantum dots (0D-QDs) of core/shell CdSe/ZnS, are converted into hybrid excitons (HXs) that are several times more weakly bound. Femtosecond pump-probe spectroscopy measurements indicated that HXs are formed through electron transfer to the 0D-QDs on a sub-45 fs time-scale. The conversion of strongly bound excitons in 2D semiconductors into easily dissociable HXs may open new opportunities for optoelectronic applications, notably in photovoltaics where exciton dissociation is of prime importance.

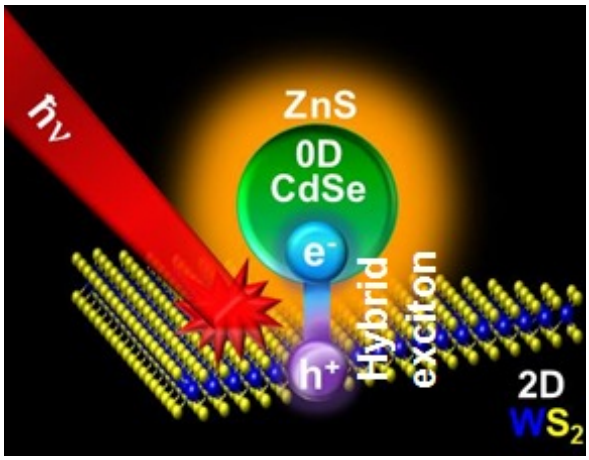

Photoexcitation of a 2D WS2 monolayer creates excitons which rapidly (sub-45 fs) dissociate and transfer electrons to 0D ZnS-capped CdSe quantum dots, as revealed by ultrafast pump-probe laser spectroscopy. The resulting 0D/2D hybrid exciton, with an electron in the quantum dot and a hole in the 2D layer, is more weakly bound than the original 2D exciton.

This work was performed at the Center for Nanophase Materials Sciences. Synthesis of 2D-WS2 materials was supported by the Materials Science and Engineering Division, Office of Basic Energy Sciences, U.S. Department of Energy.

Publication: A. Boulesbaa, Kai Wang, Masoud Mahjouri-Samani, Mengkun Tian, Alexander A. Puretzky, Ilia Ivanov, Christopher M. Rouleau, Kai Xiao, Bobby G. Sumpter, David B. Geohegan J. Am. Chem. Soc. 138 (2016) 14713-14719. DOI: 10.1021/jacs.6b08883

Science Highlights (Continued)

Theoretical Foundation for Probing the Pairing Mechanism in Iron-based Superconductors using Raman Spectroscopy

S. Maiti, T. A. Maier, T. Böhm, R. Hackl, and P. J. Hirschfeld

Researchers developed a general theory for the Raman response of multi-band systems. This theory was then used to identify new spectroscopic features in unconventional superconductors.

This work shows how Raman spectroscopy can identify multiple sub-leading pairing channels as collective modes in addition to the ground state.

In unconventional superconductors, understanding the form of the pairing interaction that gives rise to superconductivity is the primary goal. In the iron-based superconductors, the pairing appears to be s-wave, with the order parameter changing sign between the electron and hole Fermi surface pockets. The nature of the interaction that gives rise to this gap structure, however, continues to be a matter of wide debate. Theoretical calculations based on spin fluctuation theory do indeed find an s-wave ground state and predict multiple subleading d-wave states that can be strongly competitive. From previous work of Bardasis and Schrieffer it is known that a subdominant d-wave pairing channel can lead to a collective excitonic mode, a Bardasis-Schrieffer mode, that consists of two quasiparticles in a Cooper pair d-wave state. The energy of this mode lays below the pair-breaking edge of the condensed s-wave system and can be detected in electronic Raman scattering. In this work, a theory of the Raman response was developed for systems like the iron-based superconductors, in which not one, but multiple pairing channels are sub-dominant. This theory identifies the signatures of multiple Bardasis-Schrieffer type collective modes and connects the weights of these modes to the subleading gap structures predicted by the calculations. This work thus delivers an important contribution for interpreting the Raman response of multi-band superconductiors, thereby establishing Raman spectroscopy as a unique and useful tool for probing the predictions of pairing theories.

In this regard, Raman spectroscopy is a very useful tool, as it identifies the ground state and also the subleading pairing channels by probing collective modes.

This research was conducted at the Center for Nanophase Materials Sciences, which is a Department of Energy (DOE) Office of Science User Facility.

Publication: S. Maiti, T. A. Maier, T. Böhm, R. Hackl, P. J. Hirschfeld, “Probing the Pairing Interaction and Multiple Bardasis-Schrieffer Modes Using Raman Spectroscopy”, Phys. Rev. Lett. 117, 257001 (2016).

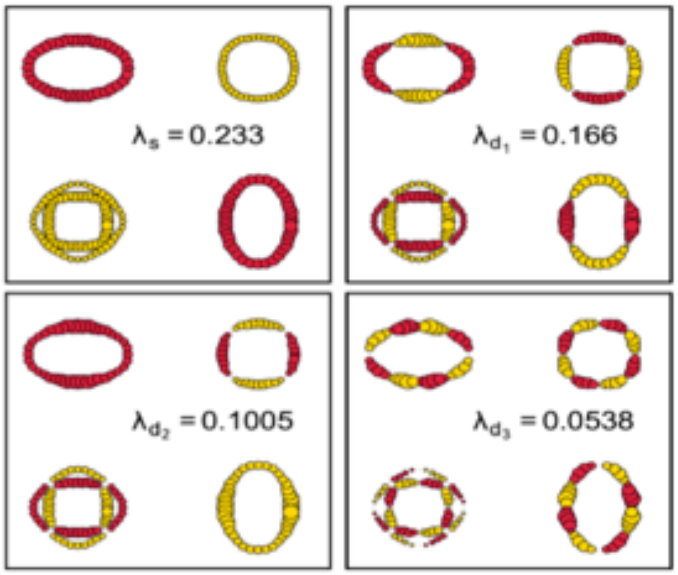

Leading s-wave (top left) and subleading d-wave gap structures calculated for hole doped BaFe2As2.

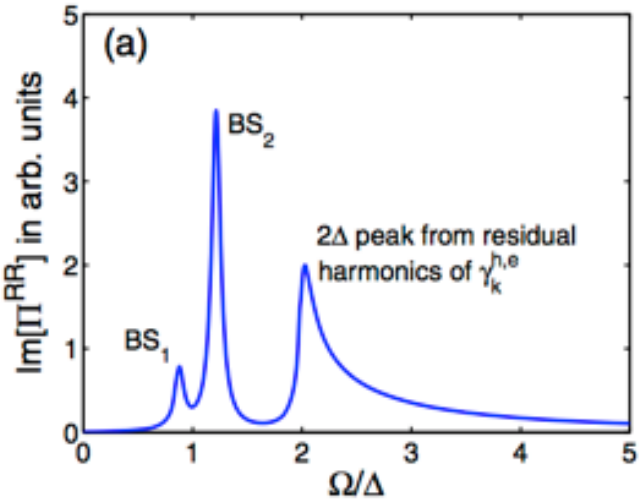

Calculated Raman response with multiple in-gap Bardasis-Schrieffer collective modes

Science Highlights (Continued)

Approach Developed for Automatically Characterizing Parallel Application Communication Patterns

P.C. Roth, J.S. Meredith, and J.S. Vetter

Researchers developed an approach for auto- matically recognizing and concisely representing the communication patterns of parallel applications that use the Message Passing Interface (MPI).

Characterizing parallel application communication patterns requires consid- erable expertise but greatly simplifies tasks such as proxy application validation, performance problem diagnosis, and debugging. Our automated approach significantly reduces the expertise necessary for effective characterization.

Funding for this work was pro- vided by the Office of Advanced Scientific Comput- ing Research, U.S. Department of Energy. The work was performed at Oak Ridge National Laboratory (ORNL).

To enable characterization of parallel application communication patterns by non-experts, we developed an approach for automatically recog- nizing and parameterizing communication patterns in MPI-based applications. Beginning with a communication matrix that indicates how much data each process transferred to every other process during the application's run, we use an automated search to recognize communication patterns within this matrix. At each search step, we recognize patterns from a pattern library in the communication matrix. Using a technique similar to astronomy's "sky subtraction," when we recognize a pattern we remove it from the matrix and apply our recognition approach recursively to the resulting matrix. Because more than one pattern might be recognized at each search step, the search produces a search results tree whose paths between root and leaves represent collections of patterns recognized in the original matrix. The path that accounts for the most of the original communication matrix's traffic corresponds to the collection of patterns that best explains the application's communication behavior. We implemented our approach in a tool called AChax that was highly effective in recognizing the communication patterns in a synthetic communication matrix and the regular communication patterns in matrices obtained from the LAMMPS molecular dynamics and LULESH shock hydrodynamics applications.

Publication: P.C. Roth, J.S. Meredith, J.S. Vet- ter, "Automated Characterization of Parallel Ap- plication Communication Patterns," 24th Inter- national ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC-2015), Portland, Oregon, June 2015.

Recognizing and removing the contribution of a 2D nearest neighbor pattern in a synthetic communication matrix. This represents one step in our automated search- based approach.

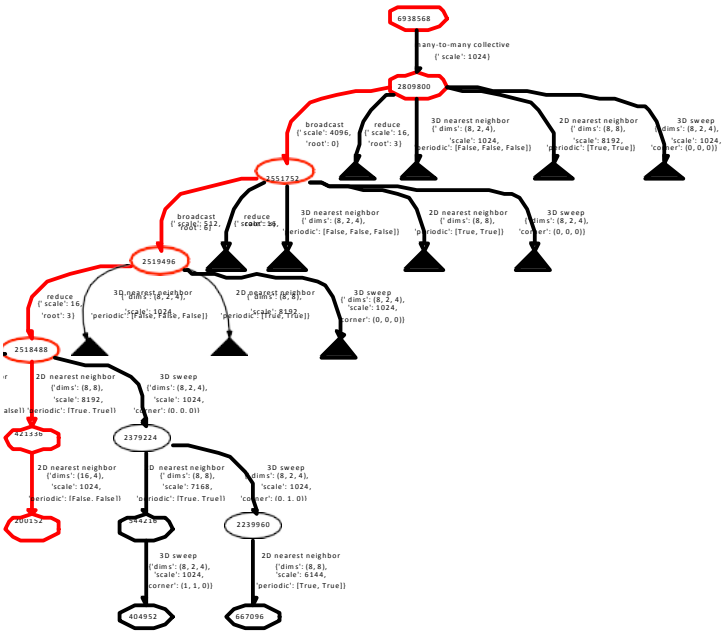

Results tree produced by search-based auto- mated pattern characterization. Nodes are labeled with the amount of communication yet to be explained, and edges are labeled with the recognized pattern. The path in red indicates the patterns that best explain the original com- munication behavior

News

Nano-spike catalysts convert carbon dioxide directly into ethanol

In a new twist to waste-to-fuel technology, scientists at the Department of Energy's Oak Ridge National Laboratory have developed an electrochemical process that uses tiny spikes of carbon and copper to turn carbon dioxide, a greenhouse gas, into ethanol. Their finding, which involves nanofabrication and catalysis science, was serendipitous.

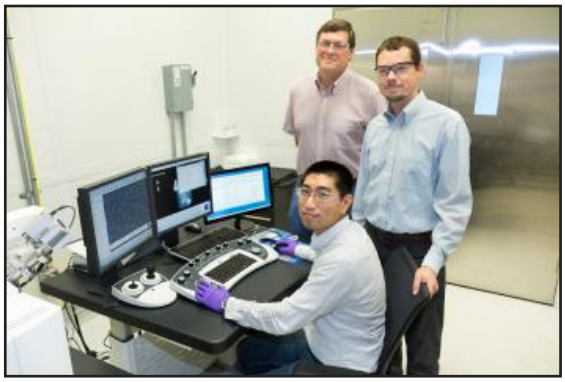

ORNL's Yang Song (seated), Dale Hensley (standing left) and Adam Rondinone examine a carbon nanospike sample with a scanning electron microscope.

"We discovered somewhat by accident that this material worked," said ORNL's Adam Rondinone, lead author of the team's study published in ChemistrySelect. "We were trying to study the first step of a proposed reaction when we realized that the catalyst was doing the entire reaction on its own."

Read the full story here - https://www.ornl.gov/news/nano-spike-catalysts-convert-carbon-dioxide-directly-ethanol.

Community Service

- Rizwan A. Ashraf, program committee, Second International Workshop on Heterogeneous High-Performance Reconfigurable Computing (H2RC)
- David E. Bernholdt, scientific committee, 1st International Conference on Computational Methods and Algorithms on HPC Platforms and Accelerators (CompHPC 2017)
- David E. Bernholdt, lead organizer, Birds of a Feather: Software Engineering for Computational Science and Engineering on Supercomputers at SC16
- Christian Engelmann, peer reviewer, IEEE Transactions on Nuclear Science
- Christian Engelmann, peer reviewer, IEEE Transactions on Parallel and Distributed Systems (TPDS)
- Saurabh Hukerikar, program committee, 25th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP 2017)
- Saurabh Hukerikar, peer reviewer, Future Generation Computer Systems (FGCS)
- Seyong Lee, program committee, IPDPS17: IEEE International Parallel & Distributed Processing Symposium
- Seyong Lee, external reviewer, TWMS: Journal of Pure and Applied Mathematics
- Thomas J. Naughton, peer reviewer, ACM Computing Surveys (CSUR)
- Phil Roth, chair, SC16 Tutorials
- Phil Roth, executive director, SC17
- Phil Roth, guest editor, Special Issue of Parallel Computing: Systems & Applications, to appear Spring 2017
- Phil Roth, reviewer, DOE SBIR Phase I
- Phil Roth, program committee, HPC Pipeline Workshop
- Phil Roth, reviewer, Cluster Computing Journal
- Geoffroy R. Vallee, peer reviewer, IEEE Transactions on Parallel and Distributed Systems (TPDS)

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).

Events (Continued)

CAM Seminar Series

- Dr. Cory Hauck, Oak Ridge National Laboratory - Discontinuous Galerkin Methods for Transport Equations and the Diffusion Limit
- Dr. Eirik Endeve, Oak Ridge National Laboratory - A Realizability-preserving DG Method for the Two-moment Model of Radiative Transfer
- Dr. Valeriy A. Buryachenko, Micromechanics & Composites - Micromechanics of Random Structural Composits with Nonlocal Thermoelastic Properties of Constituents
- Mr. Peter Jantsch, University of Tennessee, Knoxville - Lebesgue constant for Weighted Leja Sequences on Unbounded Domains
- Mr. Nicholas Dexter, Oak Ridge National Laboratory/ University of Tennessee - Global Reconstruction of Solutions to Parametric PDEs via Compressed Sensing
- Dr. Abdellah Chkifa, Oak Ridge National Laboratory - A Sparse Grid Collocation Method Based on LaVallee Poussin Kernel
- Dr. Ed D'Azevedo, Oak Ridge National Laboratory - Delayed Update Algorithms for Quantum Monte Carlo Simulations
- Dr. Minh-Binh Tran, University of Wisconsin, Madison - Quantum Kinetic Problems
- Dr. Manda Winlaw, University of Waterloo - Nonlinearly Preconditioned Conjugate Gradient Algorithm for Rank-R Canonical Tensor Approximation and Latent Factor Model Solutions
- Dr. Rick Archibald, Oak Ridge National Laboratory - Sparse Sampling Methods for Neutron Data
- Dr. Pablo Seleson, Oak Ridge National Laboratory - Convergence Studies for Nonlocal Peridynamic Models

CSMD Seminars

- Mr. Jason Wang: Data I/O middleware for isolating development complexities and enabling real-time optimization
- Graham Lopez: Moving Today's HPC Applications to Upcoming Heterogeneous Architectures
- Kshitij Mehta: Accelerating Seismic Applications for GPUs using Directive-Based Programming Models
- Dr. Nilesh Mahajan: Kanor: A Language for Declarative Communication
- David Riegner and Wolfgang Windl: Center for Performance and Design of Nuclear Waste Forms and Containers (WastePD EFRC)
- Karl Fuerlinger: DASH: A C++ PGAS Library for Distributed Data Structures and Parallel Algorithms
- Hongzhang Shan: Optimizing the nearest-neighbor communication using one-sided communication
- Kevin Lai: SOFC-MP and Recent Developments
- Ranadip Acharya: Prediction of microstructure in laser powder bed fusion process
- Harsh Bhatia: Scientific Visualization and Analysis for Data-driven Discovery

Awards and Recognition

Seyong Lee wins 2016 IEEE-CS TCHPC Award

Seyong Lee (Future Technologies Group, CSMD) has been selected as one of three recipients of the 2016 IEEE-CS TCHPC Award for Excellence for Early Career Researchers in High Performance Computing (https://www.computer.org/web/pressroom/TCHPC-Award-2016). The IEEE Computer Society TCHPC Award for Excellence for Early Career Researchers in High Performance Computing, sponsored by the IEEE Computer Society Technical Consortium on High Performance Computing (TCHPC) and its member Technical Committees, Technical Committee on Parallel Process (TCPP) and Technical Committee on Computer Communications (TCCC), recognizes up to 3 individuals who have made outstanding, influential, and potentially long-lasting contributions in the field of high performance computing within 5 years of receiving their PhD degree as of January 01 of the year of the award. The awards will be presented at the SC16 conference that will be held in Salt Lake City, UT, USA during November 13 - 18, 2016.

The IEEE Computer Society TCHPC Award for Excellence for Early Career Researchers in High Performance Computing, sponsored by the IEEE Computer Society Technical Consortium on High Performance Computing (TCHPC) and its member Technical Committees, Technical Committee on Parallel Process (TCPP) and Technical Committee on Computer Communications (TCCC), recognizes up to 3 individuals who have made outstanding, influential, and potentially long-lasting contributions in the field of high performance computing within 5 years of receiving their PhD degree as of January 01 of the year of the award.

Nageswara Rao Part of Team Recognized with Best Paper Award

CSMD researcher Nageswara Rao's paper "Experimental Analysis of File Transfer Rates Over Wide-Area Dedicated Connections" received a best paper award at the 18th IEEE International Conference on High Performance Computing and Communications. The HPCC 2016 conference is the 18th edition of a forum for engineers and scientists in academia, industry, and government to address the resulting profound challenges and to present and discuss their new ideas, research results, applications and experience on all aspects of high performance computing

Awards and Recognition (Continued)

and communications. IEEE HPCC 2016 is sponsored by IEEE, IEEE Computer Society, and IEEE Technical Committee on Scalable Computing (TCSC).

Other authors include: Qiang Liu, Satyabrata Sen, Gregory Hinkel, Neena Imam, Ian Foster, Rajkumar Kettimuthu, Bradley Settlemyer, Chase Q. Wu, and Daqing Yun.

George Ostrouchov AAAS Fellow

George Ostrouchov is a senior research staff member and analytics task lead in the Scientific Data Group of ORNL's Computer Science and Mathematics Division. He also serves as a matrix member of the Advanced Data and Workflow Group at the Oak Ridge Leadership Computing Facility and as a Joint Faculty member of the Business Analytics and Statistics Department at the University of Tennessee, Knoxville.

Ostrouchov is cited by the AAAS "for distinguished leadership in the field of statistical computing, particularly to enable parallel computation on big data with statistical software, and for service to profession." He is also a fellow of the American Statistical Association. His current research enables scalable data analysis on large systems, particularly for multivariate methods in high-dimensional data and dimension reduction. He leads the pbdR.org project for scaling the R language to large distributed computing systems.

He received a B.Math. in statistics from the University of Waterloo in Canada, followed by an M.Sc. and a doctorate in statistics from Iowa State University. Throughout his career he has worked on various "large" aspects of statistics and data analysis. He is author or co-author of more than 100 publications on topics including sparse matrix computations, model search algorithms, Bayesian estimation, streaming and parallel data algorithms, and dimension reduction, applied across diverse science domains.

Jeffrey Vetter IEEE Fellow

Jeff Vetter was elevated to IEEE Fellow for his contributions to high performance computing.

Recognizing the achievements of its members is an important part of the mission of the IEEE. Each year, following a rigorous evaluation procedure, the IEEE Fellow Committee recommends a select group of recipients for elevation to IEEE Fellow. Less than 0.1% of voting members are selected annually for this member grade elevation.

R&D100 Award

Researchers at the Department of Energy's Oak Ridge National Laboratory have received seven R&D 100 Awards in recognition of their significant advancements in science and technology. The honorees were recognized on Nov. 3 at the 54th annual R&D 100 Conference, sponsored by R&D Magazine.

The awards, known as the "Oscars of Invention," honor innovative breakthroughs in materials science, biomedicine, consumer products and more from academia, industry and government-sponsored research agencies.

One of the seven awards went to the Virtual Environment for Reactor Applications (VERA) project.

VERA represents a new generation of capabilities that can be used by the nuclear industry to improve the operation of nuclear reactors. It provides high-fidelity, science-based simulation tools to address key challenges, enabling increased clean, safe and reliable generation of electricity. The numerical algorithms and design embodied in VERA and its components provide a foundation that advances the state of the art in nuclear simulation tools, able to take advantage of both current and future computational resources ranging from industry clusters to leadership-class supercomputers. VERA's user inter-face allows easy creation of complex models, making the use of high performance computing for engineering and design more available and usable. VERA has been proven in real-world applications through the simulation of Tennessee Valley Authority's Watts Bar Nuclear Plant, first reproducing 18 years of operation for the first unit, and currently modeling startup of the second unit.

VERA is a result of the collaboration of Jess Gehin (Nuclear Science & Engineering), Kevin Clarno, Benjamin Collins, Gregory Davidson, Thomas Evans, Andrew Godfrey, Seth Johnson, Steven Hamilton, Douglas Kothe, Rose Montgomery, Robert Salko, Srdjan Simunovic, Stuart Slattery, Shane Stimpson, John Turner, Mark Baird, Tara Pandya, Kang Seog Kim, Matthew Sieger and partners from EPRI, INL, LANL, Sandia, North Carolina State University, University of Michigan, Westinghouse Nuclear and Core Physics.

New Faces

Joel E. Denny

Computer Scientist, Future Technologies Group

Please join us in welcoming Dr. Joel E. Denny, who began a new role as a Computer Scientist in the Future Technologies Group, Computer Science and Mathematics Division, on October 10, 2016.

Joel received his Ph.D. in Computer Science in 2010 from Clemson University. After his Ph.D., he worked as a complier engineer in industry. Two years ago, he joined the Joint Institute for Computational Sciences with the University of Tennessee and ORNL. Since then, he has worked on various LLVM-based compiler research projects in the Future Technologies Group. Currently, his research focuses on complier techniques for programming non-volatile memory.

Joel is located in Building 5100, Room 236, and can be reached at 865-576-7592.

Andrey Prokopenko

Andrey joins ORNL after earning a Ph.D. in Applied Mathematics from the University of Houston and then working as a postdoc in the Scalable Algorithms Department in the Center for Computing Research of the Sandia National Laboratories. His research is focused upon scalable numerical algorithms with an emphasis on linear algebra and multigrid methods. Andrey is a significant contributor to the well-known Trilinos project and is a lead developer of a MueLu multigrid library.

Shirley Moore

Senior Computer Scientist, Future Technologies Group

Please join us in welcoming Dr. Shirley Moore to the Future Technologies Group in the Computer Science and Mathematics Division at ORNL. Shirley comes to ORNL after working for four years as an Associate Professor of Computer and Computational Sciences at the University of Texas at El Paso, where she helped develop the interdisciplinary graduate Computational Science Program. Prior to that, she helped a research position with her Innovative Computing Laboratory at the University of Tennessee. She earned a Ph.D. in Computer Sciences from the Purdue University in 1990. Her research is the area of performance modeling and optimization of high performance computing systems and applications.

Shirley is located in Building 5100, Room 234, and can be reached at 576-4838.

Ferrol Aderholdt

Computer Science Research Group

Ferrol Aderholdt received his Ph.D. degree in Electrical and Computer Engineering from Tennessee Technological University under the supervision of Dr. Stephen L. Scott. While working toward his doctorate, he studied the survivability of modern cloud computing systems and proposed mechanisms and methods to increase the survivability of these systems. Ferrol's work at ORNL is focused on the abstraction of heterogeneous memories in current and future extreme scale systems.

Rizwan Ashraf

Computer Science Research Group

Rizwan Ashraf received his Ph.D. in Computer Engineering from the University of Central Florida where he focused on research and development of adaptive strategies for fault-resilient energy-aware computing at both hardware and software levels. His expertise involves the investigation and efficient mitigation of the effects of hard faults due to transistor aging, manufacturing-induced process variations, etc., and transient faults due to high-energy radiation particles in computing platforms ranging from small-scale, e.g., FPGAs, to large-scale, e.g., tracking transient faults in scientific supercomputing applications. His focus at ORNL will be to develop a structured approach for fault resilience in extreme-scale supercomputers using design pattern abstractions. He will build multi-level adaptive fault resilience architectures in hardware and software and their coordination strategies.

Ozgur Cekmer

Computer Science Research Group

Ozgur Cekmer received his Ph.D. degree in Mechanical Engineering from the University of Tennessee, Knoxville, where he worked with Professor Matthew M. Mench on development of a combined path-percolation and lattice Boltzmann model, which is used to simulate two-phase flow through inhomogeneous porous media. Ozgur's current research focuses on the behavior of helium and hydrogen gas in bulk tungsten, which is a key aspect of modeling of tokamak reactors for the production of energy from fusion, using continuum reaction-diffusion0advection cluster dynamics, and applying uncertainty quantification analysis on the developed cluster dynamics code.

New Faces (Continued)

Yawei Hui

Computer Science Research Group

Yawei Hui received his M.S. degree in Computer Science from the University of Tennessee, Knoxville in 2015. He also received his Ph.D. in Astrophysics from the Johns Hopkins University in 2008. His current research interests lie in distributed computing, scientific and high performance computing, data mining and machine learning, scientific data processing, and visualization. At ORNL, he works with many staff members from CSMD, CNMS, IFIM on developing the Bellerophon Environment for Analysis of Materials (BEAM) software package which would enable the researchers to perform complex data analysis on big data sets at scale by using the supercomputing power provided via Oak Ridge National Laboratory-based user facilities such as CADES and OLCF (Titan, Rhea, etc.).

Saurabh Hukerikar

Computer Science Research Group

Saurabh Hukerikar received his Ph.D. from the Department of Electrical Engineering at the University of Southern California. At USC, he worked with the Computational Systems and Technology Division at the Information Sciences Institute. The emphasis of his graduate work was on designing programming model and runtime solutions to address the challenge of resilience for extreme-scale high-performance computing (HPC) systems. At ORNL, Saurabh is developing resilience-driven design patterns for extreme-scale HPC systems. The effort seeks to develop mechanisms based on formalized best practices in HPC resilience. The resilience design patterns will enable HPC system designers, programmers, and users to develop coordinated cross-layer resilience strategies, and provide a structured approach to navigate the trade-off space between performance, resilience, and power consumption.

Tonglin Li

Computer Science Research Group

Tonglin Li received his Ph.D. degree in Computer Science from Illinois Institute of Technology at Chicago, where he worked with Professor Ioan Raicu on distributed storage systems and applications. At ORNL, Tonglin's work focuses on the design and implementation of integrated data storage services. These services will help large scale scientific applications to better utilize the resource of current and next generation leadership supercomputers.

William F. Godoy

Computer Scientist, Scientific Data Group (SDG)

William comes to ORNL after being a R&D staff member at Intel Corporation where he developed large-scale thermal reliability models and software products in collaboration with design and manufacturing teams for the next generation of processors. William received his Ph.D. in Mechanical Engineering from SUNY at Buffalo. He then joined NASA Langley Research Center as a postdoc. His research focused on high-performance computing techniques such as parallel hybrid Message Passing Interface/threads, general-purpose graphics processing units, and datasets reduction applied on large-scale radiative transfer problems. William joins SDG to work on the next generation of Adaptive Input/Output Systems (ADIOS) research methods for Exascale applications.

William in located in Building 4100, Office C232 and can be reached at godoywf@ornl.gov or 865-574-8878.

Ruonan (Jason) Wang

Research Associate (UTK), Scientific Data Group (SDG)

Jason is currently a research associate at UTK and is working with the Scientific Data Group at ORNL. Jason has finished his Ph.D. at the University of Western Australia, being involved in the data storage system design for the future world's largest radio telescope, the Square Kilometre Array. His research has focused on prototyping next generation data Input/Output middleware, which enables run-time optimizations. He will officially join ORNL soon after he officially receives his Ph.D.

Jason is located in Building 4100, Office C238 and can be reached at wangr1@ornl.gov or 865-574-8984.

New Faces (Continued)

Lipeng Wan

Postdoctoral Research Associate, Scientific Data Group

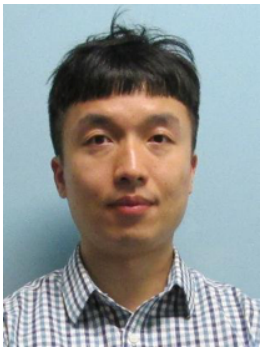

Lipeng joins ORNL after earning a Ph.D. in computer science major from the University of Tennessee, Knoxville. His research focuses on parallel and distributed storage systems, high performance computing (HPC) and Input/Output (I/O), and performance optimization for solid-state drives. He brings experience modeling the performance and reliability of large-scale HPC systems, including the I/O behaviors of large-scale scientific applications, the failure patterns and runtime characteristics of the HPC systems, etc.

Lipeng is located in Buidling 4100, office C222 and can be reached at wanl@ornl.gov or 865-574-4703.

Mark Kim

Postdoctoral Researcher, Scientific Data Group (SDG)

Mark Kim is a postdoctoral researcher in the SDG at Oak Ridge National Laboratory. He received his Ph.D. in Computing in 2015 from the University of Utah under Charles "Chuck" Hansen. His research interests include data-parallel processing for visualization, flow visualization, and large-scale visualization.

Mark is located in Building 4100, Office A251 and can be reached at kimmb@ornl.gov or 865-241-6877.

Matt Wolf

Senior Computer Scientist, Scientific Data Group

Matthew Wolf comes to Oak Ridge after being a research scientist at Georgia Tech for more than 15 years, the last 13 of which he was also a joint appointment faculty to ORNL. His background in Mathematics and Computational Physics, as well as, High Performance Computing, shapes much of his research interest in computer operating systems and middleware for large data systems. His research targets high performance, scalable applications, with particular focus on Input/Output (I/O) and adaptive streaming middlewares. Specific research topics include adaptive I/O interfaces, metadata-rich data services, creation of dynamic in situ workflows, and the handling and fusion of heterogeneous data types.

Matt is located in Building 4100, Office C226, and can be reached at wolfmd@ornl.gov or 865-576-0324.

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 nine Groups. These Groups and their Group Leaders are:

- Complex Systems – Jacob Barhen
- Computational Biomolecular Modeling & Bioinformatics – Mike Leuze
- Computational Chemical and Materials Sciences – Bobby Sumpter
- Computational Earth Sciences – Kate Evans
- Computational Engineering and Energy Sciences – John Turner
- Computational Applied Mathematics – Clayton Webster
- Computer Science Research – David Bernholdt
- 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

OAK RIDGE National Laboratory