ISSUE
20
SPRING
2016

QUARTERLY NEWSLETTER FOR THE COMPUTER SCIENCE AND MATHEMATICS DIVISION

CSMD

this issue

Science Highlights P.01
News P.07
Awards and Recognition P.08
New Staff P.08
Community Service P.08
Events P.09
Software P.11
Publications of Note P.11
About CSMD P.15

Science Highlights

Extending Supramolecular Polymerization to New Lengths

Jun Wang, Peter V. Bonnesen, E. Rangel, E. Vallejo, Ariadna Sanchez-Castillo, H. James Cleaves II, Arthur P. Baddorf, Bobby G. Sumpter, Minghu Pan, Petro Maksymovych, and Miguel Fuentes-Cabrera

A molecule, called a nucleoside analog and which is composed of an Adenine moeity and glycol group, was deposited on top of the Au(111) surface and studied with scanning tunneling microscopy and density functional theory calculations. The molecule was found to self-assemble into a supramolecular polymer, i.e., a polymer held by hydrogen bonds, that is not only two order of magnitudes longer than any other polymeric structure made by similar molecules on substrates but also encodes a sequence which is erasable with electron-induced excitation. The sequence encoded was the following: glycol-glycol-Adenine-Adenine-glycol-glycol, etc., and it remained intact for hundreds of nanometers. This work suggest that nucleoside polymerization on patterned substrates could be used to create sequence-controlled polymers that store information that is readable and erasable with microscopy techniques. Further, because polymerization takes place on a susbtrate, theory can be used to design what patterns and what molecules will produce the desired polymeric outcomes.

Publication: "*Supramolecular polymerization of a prebiotic nucleoside provides insights into the creation of sequence-controlled polymers*"; Jun Wang, Peter V. Bonnesen, E. Rangel, E. Vallejo, Ariadna Sanchez-Castillo, H. James Cleaves II, Arthur P. Baddorf, Bobby G. Sumpter, Minghu Pan, Petro Maksymovych, and Miguel Fuentes-Cabrera; Scientific Reports (2015) DOI: 10.1038/srep18891

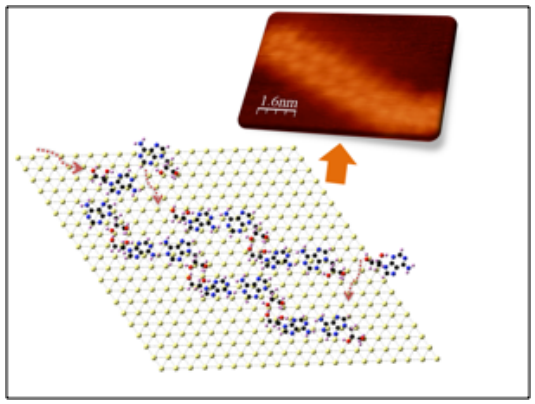

Schematic representation of the polymerization process. Upper right-corner, high-resolution image of the chain-like structure.

This research was conducted at the Center for Nanophase Materials Sciences (CNMS), which is a DOE Office of Science User Facility. The computations were performed using resources of the CNMS and the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory. We kindly thank Divina Anunciado and Hugh O'Neill for the circular dichroism measurements. One of us, A. S.-C. acknowledges the financial support from PROMEP, grant ACTA-910. Parts of the calculations were performed in the DGCTIC-UNAM supercomputing center and ESAp-UAEH.

OAK RIDGE National Laboratory

Science Highlights (Continued)

Revealing Various Stacking Patterns Between Layers in Two-dimensional (2D) Materials

Alexander A. Puretzky, Liangbo Liang, Xufan Li, Kai Xiao, Bobby G. Sumpter, Vincent Meunier, and David B. Geohegan

In this work unique twisted bilayers of MoSe2 with periodic multiple stacking configurations and interlayer couplings were discovered in the narrow range of twist angles, 60± 3°, using ulra-low frequency Raman spectroscopy and first-principle theory. We showed that the slight deviation from 60° creates patches featuring all three high-symmetry stacking configurations (2H or AA', AB', and A'B) in one unique bilayer system. In this case, the periodic arrangement of the patches and their size strongly depend on the twist angle. Our first-principle modeling predicts significant changes in frequencies and intensities of low-frequency modes versus stacking and twist angle. Experimentally, the variable stacking and coupling across the interface are revealed by the appearance of two breathing modes, corresponding to the mixture of the high-symmetry stacking configurations and unaligned regions of monolayers. Only one breathing mode is observed outside the narrow range of twist angles. This indicates a stacking transition to unaligned monolayers with mismatched atom registry without the in-plane restoring force required to generate a shear mode. The variable interlayer coupling and spacing in transition metal dichalcogenide bilayers revealed in this study may provide a new platform for optoelectronic applications of these materials.

The Raman spectroscopy part of this research, including aspects of theory, was conducted at the Center for Nanophase Materials Sciences, a U.S. Department of Energy Office of Science User Facility. The synthesis science including CVD was supported by the U.S. Department of Energy, Office of Science, Basic Energy Sciences, Materials Sciences and Engineering Division. The theoretical work at Rensselaer Polytechnic Institute (RPI) was supported by NSF EFRI-2DARE 1542707. L.L. was supported by a Eugene P. Wigner Fellowship at the Oak Ridge National Laboratory. The computations were performed using the resources of the Center for Computational Innovation at RPI.

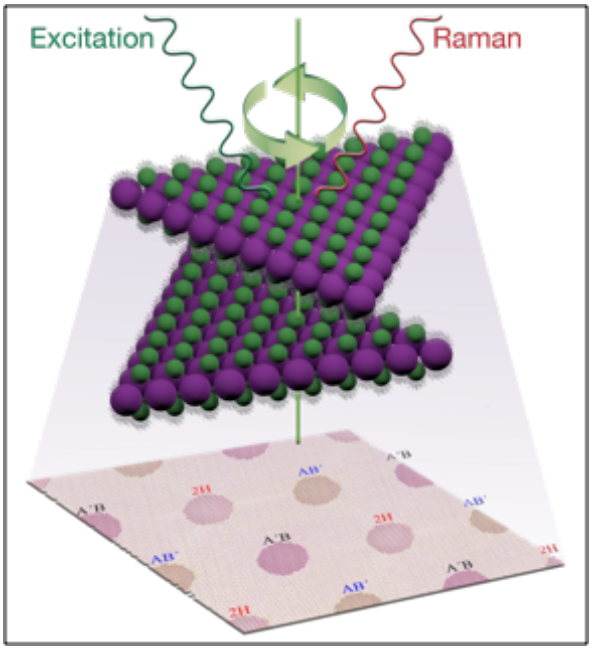

Mutual rotation of two monolayers of transition metal dichalcogenides creates large variety of bilayer stacking patterns depending on a twist angle. The most interesting patterns, periodic arrangement of the high-symmetry patches (AB', A'B, 2H), occur at twist angles near 60° as shown in the projection. Fast and efficient characterization of these stacking patterns is very important for their potential applications, for example in optoelectronics. We discovered that low-frequency Raman scattering can identify these unique bilayers much faster than expensive and time-consuming atomic resolution electron microscopy.

Publication: "*Twisted MoSe2 Bilayers with Variable Local Stacking and Interlayer Coupling Revealed by Low-Frequency Raman Spectroscopy*"; Alexander A. Puretzky, Liangbo Liang, Xufan Li, Kai Xiao, Bobby G. Sumpter, Vincent Meunier, and David B. Geohegan; ACS Nano (2015). DOI:10.1021/acsnano.5b07807

ORNL Paper Surveys Approximate Computing Techniques

Sparsh Mittal

As rising performance demands confront plateauing resource budgets, approximate computing (AC) has become not merely attractive, but imperative. We have performed a comprehensive survey of techniques for AC in CPU, GPU and FPGA and various processor components (e.g. cache, main memory, secondary storage), along with approximate storage in SRAM, DRAM/eDRAM, and non-volatile memories, e.g. Flash, STT-RAM, etc. We review design of approximate circuits such as adder, multiplier and general logic circuits. We also review strategies for finding approximable program portions and monitoring output quality, and programming frameworks for AC. Finally, we have highlighted the challenges that are worthy of further investigation.

Publication: Sparsh Mittal, "A Survey of Techniques for Approximate Computing", accepted in ACM Computing Surveys 2016

This work was performed at Oak Ridge National Laboratory.

Science Highlights (Continued)

Layout-Aware Data Scheduler (LADS) for Bulk Data Transfer over Terabit Networks

Y. Kim, S. Atchley, G. Vallee, S. Lee, and G. Shipman

Researchers developed an end-to-end data transfer framework, named LADS, for bulk data transfer for terabit networks that is optimized for parallel file systems on both the source and sink. This work includes techniques to improve data transfer under congestion, without impacting other users of the parallel file systems.

This work enhances our ability to transfer bulk data between parallel file systems from different institutions over long-distance by using terabit networks, while minimizing the impact on other users of the parallel file system.

To achive this, researchers had to:

- Develop a layout-aware data scheduler for bulk data transfer that coordinates accesses to the parallel file system and the terabit network.
- Leverage high-performance capabilities of modern networks (e.g., operating system by-pass).
- Develop advanced buffering mechanisms to minimize the effect of congestion of the parallel file system.

This work was sponsored by the Department of Energy's ASCR office.

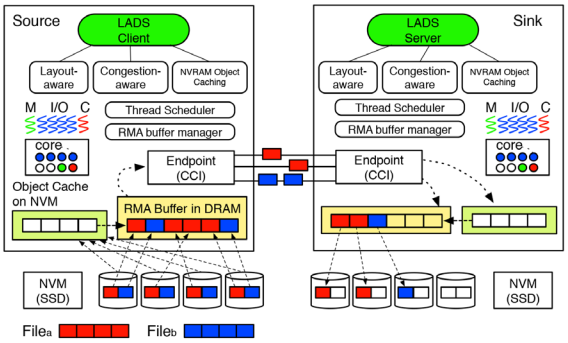

The figure presents the architecture of the LADS framework, which is based on a threaded architecture for the coordination of the accesses to the parallel file system and the network resources in order to deliver efficient end-to-end data transfer: I/O threads are in charge of accessing the parallel file system to read/write data to/from a communication buffer, while a communication thread ensures that the data is transferred efficiently over the network using the communication buffer. This architecture ensures that the overall performance of the system is not limited by the slowest components (e.g., an I/O thread can only access the file system at a slower pace because of ongoing contention) but rather benefits from the fastest ones (e.g., assigning more work to them).

Publication: Youngjae Kim, Atchley Scott, Geoffroy Vallee, Sangkeun Lee, Galen Shipman. Optimizing End-to-End Big Data Transfers over Terabits Network Infrastructure. In IEEE Transactions on Parallel and Distributed Systems (TPDS). 2016.

Data Analytics for Additive Manufacturing

R. Dehoff (PI), R. Patton, C. Steed, D. Kistler, V. Paquit and S. Powers

The large data quantities generated during Additive Manufacturing (AM) builds naturally lend themselves to the use of analytics. The work indicates that multivariate statistics can aid in the identification of areas of concern within a part, leading to potentially more robust build processes and cost saving measures.

Recent advances and success stories in AM or 3D printing have resulted in references to the process as the “next industrial revolution.” Producing parts without defects or that will adhere to failure standards are part of the challenge for the continued rapid growth of the industry. The variation in parameters during a build can contribute to the formation of defects. Most research approaches the problem from a material science perspective, but the large data quantities generated during AM builds naturally lend themselves to the use of analytics. We perform an initial investigation using a multi-pronged approach comprised of analytics and statistics for data discovery to identify areas for process improvement and promote the potential for advanced defect detection.

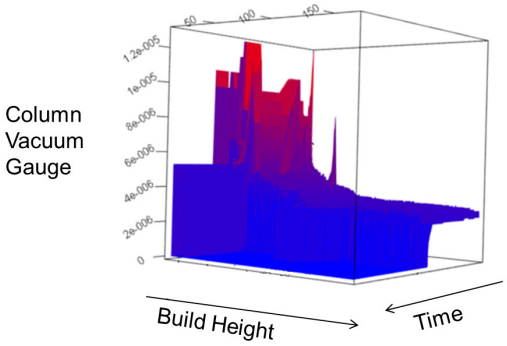

Surface plot of a parameter over time and build height, identifying statistical anomalies and regions of concern for further data analysis.

This work was funded by DOE/MDF, and was performed at ORNL.

Poster: Powers, S., Patton, R., Steed, C., Kistler, D., Paquit V. and Dehoff, R. "Application of data analytics to additive manufacturing," Conference on Data Analytics (CoDA), Santa Fe, NM, 2-4 March 2016.

Science Highlights (Continued)

Developing Fault Catalog for OLCF Systems (Titan, Jaguar, EOS)

Saurabh Gupta (CSMD-ORNL), Devesh Tiwari (NCCS-ORNL), and Christian Engelmann (CSMD-ORNL)

Researchers developed analysis tools for analyzing system logs from Titan, Jaguar and Eos systems from OLCF to extract characteristics of interest and created fault catalogue.

Understanding the characteristics of system failures in large-scale supercomputers is essential for designing mitigation techniques. It offers the research community insights into the behavior of the system. This research is developing a suite of analysis methods and gaining insights into the system errors, faults and failures. The result will be a compilation of fault taxonomy, and models that capture the observed behavior of the current systems.

To achieve this, researchers:

- Developed analysis tools for analyzing system logs from Titan, Jaguar and Eos systems.
- Combined information from different log files and created a consistent log format for analysis.
- Used previously known standard methods and some new methods to model the temporal and spatial behavior of these failure events.
- Looked at how temporal and spatial behavior has evolved over the years for these systems.
- Analyzed what is correlation of these failure types with each other (i.e., is one failure type very likely to follow another type?).
- Compared and contrasted the three systems in their behavior based on this analysis.

This work was performed at Oak Ridge National Laboratory and was sponsored by the US Department of Energy's Office of Advanced Scientific Computing Research.

Mini-Ckpts: Surviving OS Failures in Persistent Memory

David Fiala and Frank Mueller – North Carolina State University, Kurt Ferreira – Sandia National Laboratories, and Christian Engelmann – CSMD, Oak Ridge National Laboratory

Researchers developed a novel warm-reboot capability for the operating system (OS) in extreme-scale high-performance computing (HPC) systems to enable recovery from OS failures with minimal impact on running scientific applications.

The mini-ckpts framework offers a significantly more efficient solution to recover from operating system failures than traditional checkpoint/restart. It improves the efficiency and productivity of DOE's extreme-scale HPC systems.

To achieve this, researchers:

- Developed a capability to make scientific applications transparently memory resident despite OS warm reboots
- Developed a facility that enables continued operation of scientific applications after an OS warm reboot

This work was performed at North Carolina State University, Sandia National Laboratories, and Oak Ridge National Laboratory. This work was sponsored by the US Department of Energy's Office of Advanced Scientific Computing Research.

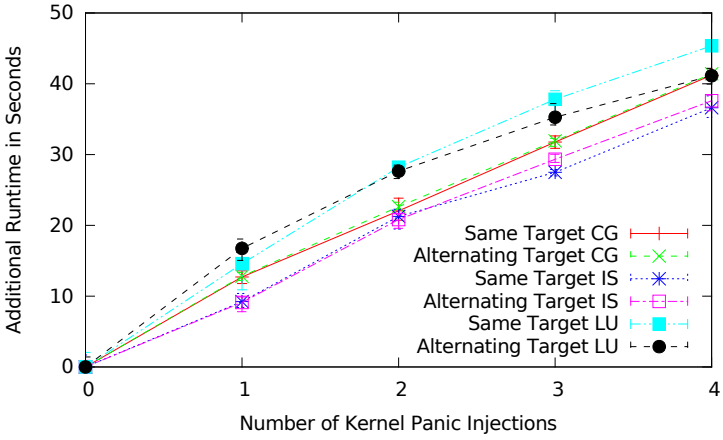

The scientific application runtime overhead stemming from the mini-ckpts recovery and rejuvenation process scales linear with the number of injected operating system failures (kernel panics).

Publication: David Fiala, Frank Mueller, Kurt Ferreira, and Christian Engelmann. Mini-Ckpts: Surviving OS Failures in Persistent Memory. In Proceedings of the 30th ACM International Conference on Supercomputing (ICS) 2016, Istanbul, Turkey, June 1-3, 2016. ACM Press, New York, NY, USA.

Science Highlights (Continued)

Theoretical Study of the Initial Stages of Self-Assembly of a Carboxysome's Facet

J. P. Mahalik, Kirsten A. Brown, Xiaolin Cheng, and Miguel Fuentes-Cabrera

Researchers devised a novel and accurate computational technique for investigating the self-assembly of large macromolecules, and used this method to reveal the initial stages of self-assembly of the carboxysome, the prototype bacterial micro-compartment. The understanding of the self-assembly of bacterial micro-compartments can help in designing artificial nano-reactors and membranes.

Bacterial microcompartments, BMCs, are organelles that exist within certain type of bacteria and act as nano-factories. Among the different types of known BMCs, the carboxysome has been studied the most. The carboxysome plays an important role in the light-independent part of the photosynthesis process, where its icosahedral-like proteinaceous shell acts as a membrane that controls the transport of metabolites. Although a structural model exists for the carboxysome shell, it remains largely unknown how the shell proteins self-assemble. Understanding the self-assembly process can provide insights into how the shell affects the carboxysome's function and how it can be modified to create new functionalities, such as artificial nanoreactors and artificial protein membranes. Here, we describe a theoretical framework that employs Monte Carlo simulations with a coarse grain potential that reproduces well the atomistic potential of mean force; employing this framework, we are able to capture the initial stages of the 2D self-assembly of

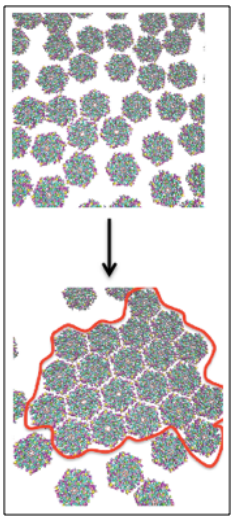

Initial and final stage of self-assembly of CcmK2 proteins

CcmK2 hexamers, a major protein-shell component of the carboxysome's facet.

The simulations reveal that CcmK2 hexamers self-assemble into clusters that resemble what was seen experimentally in 2D layers. Further analysis of the simulation results suggests that the 2D self-assembly of carboxysome's facets is driven by a nucleation-growth process, which in turn could play an important role in the hierarchical self-assembly of BMCs shells in general.

This research was conducted at the Center for Nanophase Materials Sciences (CNMS), which is a U.S. Department of Energy Office of Science User Facility. K. Brown was supported by an appointment under the Science Undergraduate Laboratory Internships (SULI), administered by the Oak Ridge Institute for Science and Education between the U.S. Department of Energy and Oak Ridge Associated Universities. The computations were performed using resources of the CNMS and the Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory.

Publication: "*Theoretical Study of the Initial Stages of Self-Assembly of a Carboxysome's Facet* " J. P. Mahalik, Kirsten A. Brown, Xiaolin Cheng, and Miguel Fuentes-Cabrera

Tailoring Plastics to Soak up Carbon Dioxide

Balaka Barkakaty, Katie L. Browning, Bobby Sumpter, David Uhrig, Ivana Karpisova, Kevin W. Harman, Ilia Ivanov, Dale K. Hensley, Jamie M. Messman, S. Michael Kilbey II, Bradley S. Lokitz

Researchers designed a polymer for selective and reversible carbon dioxide (CO2) capture.

This new polymer (based on amidines) can provide a more efficient alternative to conventional polyethyleneimine (PEI) based solid-sorbents for CO2 capture and separation.

This research was conducted at the Center for Nanophase Materials Sciences, which is a DOE Office of Science User Facility. We also thank European Commission under 7 FP (Project No. PIRSES-GA-2010-269182) and CNMS user project (Proposal No. CNMS2014-033) for sponsoring Ivana Karpisova from Comenius University, Slovakia to contribute to this work. BGS acknowledges partial support from the Center for Understanding and Control of Acid Gas-Induced Evolution of Materials for Energy (UNCAGE-ME), an Energy Frontier Research Center funded by U.S. Department of

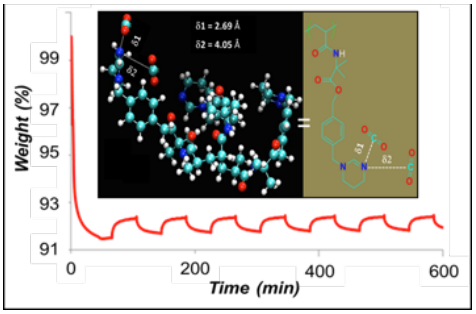

Thermogravimetric analysis (TGA) and Density functional theory (DFT) calculations coupled with modeling and simulation reveal the presence of two CO2 binding sites in the PBA-functionalized PVDMA polymer resulting in a 2-step CO2 release at room temperature.

Energy, Office of Science, Basic Energy Sciences. S.M.K acknowledges support from the National Science Foundation (Award # 1133320).

Publication: "*Amidine-Functionalized Poly(2-vinyl-4,4-dimethylazlactone) for Selective and Efficient CO2 Fixing*" Balaka Barkakaty, Katie L. Browning, Bobby Sumpter, David Uhrig, Ivana Karpisova, Kevin W. Harman, Ilia Ivanov, Dale K. Hensley, Jamie M. Messman, S. Michael Kilbey II, Bradley S. Lokitz; Macromolecules (2016) DOI: 10.1021/acs.macromol.5b02483.

Science Highlights (Continued)

A Novel Uncertainty Quantification Paradigm that Harnesses Extreme-scale Computing

Clayton Webster, Guannan Zhang, Miroslav Stoyanov, Hoang Tran and Max Gunzburger

ORNL researchers have developed a modern mathematical and statistical foundation for realizing the potential of future computing platforms, including exascale, and will ultimately enable scientists to address a fundamental question, namely "*how do the uncertainties ubiquitous in all modeling efforts affect our predictions and understanding of complex phenomena?*"

We developed a transformational methodology related to the efficient, accurate and robust computation of statistical quantities of interest (QoIs), i.e., the information used by engineers and decision makers, that are determined from solutions of complex stochastic simulations. Our scientific goal of this project is to effect a transition from currently used naïve algorithms and embedded implementations to a truly architecture-aware, predictive capability for applications that dominate the focus of the DOE mission. Examples include enhancement of reliability of smart energy grids, development of renewable energy technologies, vulnerability analysis of water and power supplies, climate change estimation, safe and cost-effective designs of current and future nuclear reactors, and accelerating the design and discovery of new materials. Guided by this grand challenge in uncertainty quantification (UQ), we constructed a rigorous mathematical framework for exploiting multicore extreme parallelism, which has be made available through ORNL Toolkit for Adaptive Stochastic Modeling and Non-Intrusive ApproximatioN (TASMANIAN), and has been demonstrated by quantifying a moderately large number of uncertainties in several DOE mission science areas listed above.

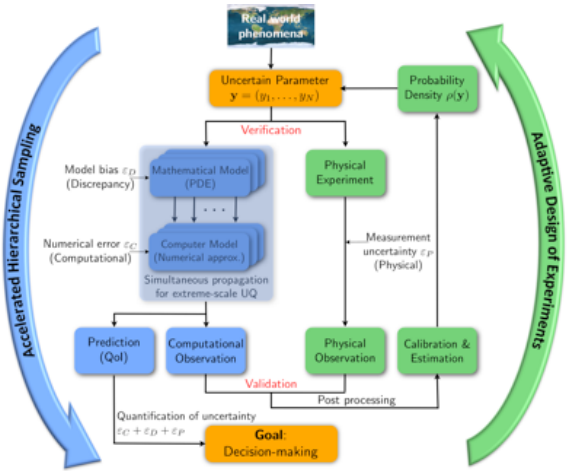

Overview of how ORNL assimilates predictions from computational simulations and experimental data at the extreme scale.

This work was performed at Oak Ridge National Laboratory and was sponsored by the US Department of Energy's Office of Advanced Scientific Computing Research.

Publications: M. Gunzburger, C. Webster, and G. Zhang, "Stochastic finite element methods for partial differential equations with random input data", Acta Numerica, 2014 vol. 23 pp. 521-650.

G. Zhang, C. Webster, M. Gunzburger, and J. Burkardt, "A Hyperspherical Adaptive Sparse-Grid Method for High-Dimensional Discontinuity Detection," SIAM Journal on Numerical Analysis vol. 53 (3) pp. 1508-1536,

M. Gunzburger and C. Webster , "An algorithmic introduction to numerical methods for uncertainty quantification for PDEs with random inputs," Springer, Applied Mathematical Sciences, 2016.

For more information about "TASMANIAN," please go to - https://tasmanian.ornl.gov.

News

Nageswara Rao's Work Recognized with Best Paper Award (March 16, 2016)

Nageswara Rao's paper "Performance Comparison of SDN Solutions for Switching Dedicated Long-Haul Connections" received a best paper award (one of two) at The International Symposium on Advances in Software Defined Networking and Network Functions Virtualization, SOFTNETWORKING 2016, February 21 - 25 in Lisbon, Portugal.

Read the full citation [here].

Researchers Develop Way to Better Predict Disease-Causing Mutations in Human Genes

Two UT researchers have developed a method that could help clinicians and scientists better predict which mutations in people's genes could cause a disease and which would remain dormant.

Ogun Adebali

Using a computational approach, Ogun Adebali and Igor Jouline have created a way to trace the evolutionary history of a human gene much more precisely.

Data gleaned from this method will help doctors know how to more effectively treat patients. It also could be used in the development of drugs to correct some of these mutations.

Igor Jouline

The study was published today in the journal Genetics in Medicinefrom the Nature Publishing Group.

Adebali, the paper's lead author, is a postdoctoral researcher and Jouline is a joint faculty professor at the UT-Oak Ridge National Laboratory's Joint Institute for Computational Sciences. They're both based in the UT Department of Microbiology.

"We're not a medical school, but this research shows that students who come to UT can tackle important medical problems," Jouline said.

Adebali and Jouline teamed up with two doctors—a physician geneticist and a specialist in lysosomal disorders. Together, they studied the Niemann-Pick disease type C, a rare inherited disorder in children that affects the way cholesterol is transported in the body.

Adebali noted that the disease can cause the spleen and liver to enlarge two to seven times their normal size. Cholesterol can also accumulate in the brain and affect speech and cognitive abilities. Most children with the disease don't survive beyond age ten. It affects one in 100,000 people in the world. For those of Jewish ancestry, the odds increase to one in 40,000.

Computational prediction of the mutation's role in disease is not a novel approach, Jouline said. What's new is the way scientists can improve the accuracy of predictions by tracing the behavior of mutations in the course of gene evolution. Adebali developed a website that will allow clinicians dealing with the Niemann-Pick disease to assess whether mutations seen in patients could be disease-causing.

"Using the same idea, we could potentially improve the diagnostics of other genetic diseases," Jouline said.

Simunovic part of team awarded patent by USPTO

Srdjan Simunovic (along with Donald L. Erdman III, Vlastimil Kunc, and Yanli Wang) was awarded patent number 9,239,277 by the USPTO. The patent, titled "Material mechanical characterization method for multiple strains and strain rates" pertains to systems and methods for characterizing mechanical properties when subject to multiple strains and strain rates.

This invention was made with government support under Contract No. DE-AC05-00OR22725 awarded by the U.S. Department of Energy.

Read the full citation [here].

Process variation threatens to slow down and even pause chip miniaturization

For past several decades, the processor industry has enjoyed the benefits of chip miniaturization and the exponential increase in the number of on-chip transistors as predicted by Moore's law. However, as process technology scales to small feature sizes, precise control of fabrication processes has become increasingly difficult. As a result, 'process variation' (PV), which refers to the deviation in parameters from their nominal specifications, has greatly worsened.

In his paper "A Survey Of Architectural Techniques for Managing Process Variation" Sparsh Mittal investigates the impact of PV along with strategies for mitigating it in a wide range of system architectures, e.g. in CPUs, GPUs, in processor components (cache, main memory, processor core), in memory technologies (SRAM, DRAM, eDRAM, non-volatile memories e.g. PCM, resistive RAM) and in both 2D and 3D processors.

Read more at: http://phys.org/news/2016-01-variation-threatens-chip-miniaturization.html#jCp

Awards and Recognition

CSMD Researcher Recognized (February 24, 2016)

Dr. Satyabrata Sen was elevated to the grade of IEEE Senior member.

Dr. Sen is a research scientist in the Center for Engineering Systems Advanced Research Group and is working on the statistical detection and estimation algorithms for radar signal processing and network fusion techniques. His research interests are in the areas of statistical signal processing, compressive sensing, asynchronous distributed tracking, network fusion, and their applications in radar, communications, and sensor arrays. He has co-authored 33 research articles in the top IEEE journals and international conferences, and contributed to 2 book chapters.

Dr. Sen is the recipient of the 2016 Sidney D. Drell Academic Award by the Intelligence and National Security Alliance (INSA) for his contributions to the U.S. Intelligence Community, and the Eugene P. Wigner Fellowship at ORNL from February 2011 to January 2013 for being an outstanding early-career professional.

New Staff

Gregory Watson

Dr. Gregory Watson is a Senior Research Scientist in the Computer Science Research Group of the Computer Science and Mathematics Division. Dr. Watson received his Bachelor's degree in Information Technology from the University of Tasmania in 1986, and his Ph. D. in Computer Science from Monash University in 2000, under the advisement of Prof. David Abramson. Between his Bachelor's degree and Ph.D., Dr. Watson held positions as Manager of Corporate Computer Systems at Griffith University, and Manager of the Queensland Supercomputing Facility, both in Brisbane, Australia. On completing his Ph.D., Dr. Watson joined the Advanced Computing Laboratory at Los Alamos National Laboratory where he contributed a number of innovations to the system software stack for state-of-the-art cluster computing systems. During this time, Dr. Watson also founded the Eclipse Parallel Tools Platform, a project he has led over the course of the last 12 years, and that continues to be used across laboratories, academia, and industry.

More recently, Dr. Watson joined IBM Research where he participated in Phase 3 of the DARPA HPCS program, and was subsequently appointed as HPC Tools Architect for IBM. He currently holds a position as Adjunct Professor in the Center for Data Science at New York University. His research interests include software engineering, tools for computational science, program debugging, distributed computing, and operating systems.

Jacek Jakowski

Jacek Jakowski started in March as a member of the Computational Chemistry & Materials Science Group

Jacek's research focuses on the development of efficient quantum dynamics methods for the application in energy research and modeling nanoscale advanced materials. Using a combination of complementary methods based on the theory of quantum mechanics, our group developed a new computational capability to study the dynamics of possible energy and materials applications. Using the NICS-managed Kraken supercomputer to perform massively parallel modeling and simulation studies in which a graphene (carbon) flake was bombarded by an ensemble of 1,000 hydrogen atoms to examine the likelihood of accumulation (adsorption) of the hydrogen on the porous graphene surface, the group found that quantum nuclear effects are responsible for increased adsorption selectivity of deuterium over hydrogen on graphene. Understanding how materials with light nuclei—most especially hydrogen—behave on nanoporous surfaces is essential in the design of new materials for energy science applications.

Community Service

- Michael Brim, peer reviewer, Journal of Parallel Computing special issue on Trends in Cluster Computing
- Michael Brim, peer reviewer, Journal of Parallel Computing special issue on DISCS 2015
- Christian Engelmann, program committee, 4th Workshop on Solving Problems with Uncertainties at the 16th International Conference on Computational Science (ICCS)
- Christian Engelmann, program committee, 6th International Workshop on Fault-Tolerance for HPC at Extreme Scale (FTXS)
- Kate Evans, presenter, Pellissippi State Computer

Community Service (Continued)

Science students tour of ORNL

- Kate Evans, panelist, Georgia Tech Society for Women Engineers' "What is the value of a Graduate Degree"
- Kate Evans, session chair, BER Exascale Requirements Workshop.
- Saurabh Hukerikar, program committee, 6th International Workshop on Fault Tolerance for HPC at eXtreme Scale (FTXS)
- Joseph Kennedy, guest speaker - League of Women Voters of Oak Ridge: Lunch with the League
- Jungwon Kim, reviewer - ACM Transactions on Architecture and Code Optimization
- Seyong Lee, reviewer - ICS (International Conference on Supercomputing) 2016 External Reviewer
- Seyong Lee, program committee - AsHES (International Workshop on Accelerators and Hybrid Exascale Systems) 2016
- Seyong Lee, external reviewer - CyS (International Journal of Computing Science and Applications)
- Sparsh Mittal, reviewer - Concurrency and Computation (Practice and Experience)
- Sarat Sreepathi, reviewer - IEEE Transactions on Evolutionary Computation
- Jeffrey Vetter, program committee - ISC 2016 Tech
- Jeffrey Vetter, program committee - ACM ICS 2016 Tech Papers
- Jeffrey Vetter, appointed associate editor - ACM Journal on Emerging Technologies in Computing (JETC)
- Geoffroy Vallee, peer reviewer, U.S. Department of Energy ASCR SBIR/STTR Phase II
- Geoffroy Vallee, strategic orientation committee, OCCIware project
- Pat Worley, session chair, BER Exascale Requirements Workshop

Events

UPCOMING

First International Workshop on Performance Portable Programming Models for Accelerators (P^3MA) - June 23, 2016

Graham Lopez

This workshop will be held on June 23, 2016, colocated with the ISC High Performance Conference in Frankfurt, Germany (http://www.isc-hpc.com).

The workshop will provide a forum for bringing together researchers, vendors, users and developers to brainstorm aspects of heterogeneous computing and its various tools and techniques.

Workshop site here - http://www.csm.ornl.gov/workshops/p3ma2016/

PAST

2nd International Workshop on the Lustre Ecosystem: Enhancing Lustre Support for Diverse Workloads

The Lustre parallel file system has been widely adopted by high-performance computing (HPC) centers as an effective system for managing large-scale storage resources. Lustre achieves unprecedented aggregate performance by parallelizing I/O over file system clients and storage targets at extreme scales. Today, 7 out of 10 fastest supercomputers in the world use Lustre for high-performance storage. To date, Lustre development has focused on improving the performance and scalability of large-scale scientific workloads. In particular, large-scale checkpoint storage and retrieval, which is characterized by bursty I/O from coordinated parallel clients, has been the primary driver of Lustre development over the last decade. With the advent of extreme scale computing and Big Data computing, many HPC centers are seeing increased user interest in running diverse workloads that place new demands on Lustre.

In early March, the 2nd International Workshop on the Lustre Ecosystem was held in Hanover, Maryland. This workshop series is intended to help explore improvements in the performance and flexibility of Lustre for supporting diverse application workloads. The first workshop, held last year in March, helped initiate a discussion on the open challenges associated with enhancing Lustre for diverse applications, the technological advances necessary, and the associated impacts to the Lustre ecosystem. This year's workshop was the second edition, and the goal was to explore The workshop program featured a day of tutorials and a day of technical paper and invited presentations, with keynote presentations leading off each day.

Presentation content for the keynote, tutorials, and technical presentations can be found at http://lustre.ornl.gov/ecosystem/agenda.html.

The workshop was organized by ORNL and sponsored by the Computational Research and Development Programs at ORNL, funded by the US Department of Defense. The workshop program co-chairs were Neena Imam, Michael Brim, and Sarp Oral. Richard Mohr of the University of Tennessee (NICS) served as the tutorials chair.

ORNL Workshop Participants:

Michael Brim (CSMD), Blake Caldwell (OLCF), Jesse Hanley (OLCF), Neena Imam (CRDP), Sarp Oral (OLCF), Nagi Rao (CSMD), and Corliss Thompson (CRDP).

Events (Continued)

CSMD Seminars

- Robert Hoy: Effect of Chain Stiffness on the Competition between Crystallization and Glass-Formation in Model Polymers
- Song Jiang: Design and Implementation of Effective Key-Value Systems for Large-scale Data Centers
- Damien LeBrun-Grandié: Modeling of the electrical, electrochemical, and thermal processes in supercapacitors
- Patrick Bridges: Integrating Performance Modeling into Scalable System Software Design
- Samer Al-Kiswany: The Old Systems and the Sea (of Applications and Hardware Changes)
- Eric Suchyta: Digging Deeper (and More Greedily) in Imaging Surveys
- Weikuan Yu: Orchestrating a Cache-Memory Concert for Massive Parallelism
- Geoffroy Vallee: Application-level fault tolerance using proposal from the MPI Forum's Fault Tolerance Working Group
- Daniel Steingart: Negotiating with Batteries
- Jacek Jakowski: Quantum Methods for Temporal and Spatial Multiphysics of Nanomaterials
- Lipeng Wan: Achieving High Reliability and Efficiency in Maintaining Large-scale Storage Systems through Optimal Resource Provisioning and Data Placement
- Mauricio Gutierrez: Efficiently simulable approximations to realistic incoherent and coherent errors and their applicability in threshold estimation
- Larry Brown: GPU Computing Update and Roadmap
- Travis Humble: CSMD Brown Bag Seminar: Quantum Computing Systems and Software
- Eric Lingerfelt: BEAM: A Computational Workflow System Enabling Scalable In Silico & Empirical Exploration of Materials Science Data in the DOE HPC Cloud
- Rizwan A. Ashraf: A Framework to Analyze the Propagation of Transient Faults in HPC Applications

CAM Seminars

- Ming-Tse Paul Laiu, University of Maryland: Positive filtered P_N moment closures for linear transport equations
- Paul Tranquilli, Virginia Tech: Lightly-Implicit Methods for the Time Integration of Extreme-Scale Applications
- Phani Motamarri, University of Michigan: Large-scale real space electronic-structure calculations using adaptive finite-element discretization
- Xu Zhang, Purdue University: Immersed Finite Element Methods for Interface Problems
- Sebastian Acosta, Baylor College of Medicine: A mathematical model of microcirculation and applications to computational hemodynamics
- Lin Mu, Oak Ridge National Laboratory: Geometric and Topological Analysis for Molecular Rigidity Functions
- Kris Garrett, Los Alamos National Laboratory: Implicitly Solving the Vlasov-Poisson System using a Phase Space Splitting
- Junping Wang, National Science Foundation: Basics of Weak Galerkin Finite Element Methods: Theory and Implementation
- Olivier Zahm, MIT: Model Order Reduction Methods for Parameter-dependent Equations - Applications in Uncertainty Quantification
- Zheng Chen, Oak Ridge National Laboratory: Recovering Exponential Accuracy in Spectral Methods Involving Piecewise Smooth Functions
- Ben Mansour Dia, King Abdullah University of Sciences and Technology: Recovering Exponential Accuracy in Spectral Methods Involving Piecewise Smooth Functions
- Nicholas Dexter, Oak Ridge National Laboratory/ University of Tennessee: Explicit Cost Bounds of Stochastic Galerkin Approximations for Parameterized PDEs with Random Coefficients
- Peter Jantsch, University of Tennessee: Analysis and Approximation of Parametric Hyperbolic PDE

Software

Lustre101 Updates

Mike Brim

The Lustre 101 web-based course series focuses on administration and monitoring of large-scale deployments of the Lustre parallel file system. Course content is drawn from large-scale Lustre experience at OLCF and other sites. Recently, a new course, "Lustre Administration Tutorials", was added to Lustre 101. This course contains full tutorials on specific aspects of Lustre administration and monitoring. Three tutorial lessons are currently available:

(1) Installing, Tuning, and Monitoring a ZFS based Lustre file system

(2) Robust Health Monitoring

(3) Network Contention and Congestion Control

Course URL - http://lustre.ornl.gov/lustre101-courses/C2/AdminTutorials.htm

Software (Continued)

OpenSHMEM 1.3 Specification

Manjunath Gorentla Venkata, Pavel Shamis, Graham Lopez, Tiffany Mintz, Mathew Baker, Aaron Welch

OpenSHMEM is a Partioned Global Address Space (PGAS) library interface specification. The goal of the specification development is to standardize the existing PGAS functionality and add functionality to the programming model to cater to the needs of OpenSHMEM users for current and emerging architectures. The OpenSHMEM specification version 1.3 is a continued evolution of the OpenSHMEM programming model towards this goal. This specification introduced various new functionalities.

Some of the highlights of functionality introduced in the specification:

- Introduced non-blocking Remote Memory Access (RMA) operations.
- Introduced atomic elemental read and write operations
- Introduced new collective operations - SHMEM_ALLTOALL and SHMEM_ALLTOALLS
- Important errata included clarifying semantics of Fence, Processing Element, and Symmetric Memory Allocation.

The ORNL team led the specification development with input from the community. Members from over 12 institutions participated in the development of the specification, which included participants from the industry, academia, and laboratories.

Work was performed with support from ORNL, ESSC, and DoD.

The OpenSHMEM specification 1.3 is available at http://openshmem.org.

Publications of Note

A Survey Of Techniques for Approximate Computing

Sparsh Mittal, ACM Computing Surveys, 2016. http://dx.doi.org/10.1145/2893356

Abstract: Approximate computing trades off computation quality with the effort expended and as rising performance demands confront with plateauing resource budgets, approximate computing has become, not merely attractive, but even imperative. In this paper, we present a survey of techniques for approximate computing (AC). We discuss strategies for finding approximable program portions and monitoring output quality, techniques for using AC in different processing units (e.g., CPU, GPU and FPGA), processor components, memory technologies etc., and programming frameworks for AC. We classify these techniques based on several key characteristics to emphasize their similarities and differences. The aim of this paper is to provide insights to researchers into working of AC techniques and inspire more efforts in this area to make AC the mainstream computing approach in future systems

A Survey of Recent Prefetching Techniques for Processor Caches

Sparsh Mittal, ACM Computing Surveys, 2016.

Abstract: As the trends of process scaling make memory system even more crucial bottleneck, the importance of latency hiding techniques such as prefetching grows further. However, naively using prefetching can harm performance and energy efficiency and hence, several factors and parameters need to be taken into account to fully realize its potential. In this paper, we survey several recent techniques that aim to improve implementation and effectiveness of prefetching. We characterize the techniques on several parameters to highlight their similarities and differences. The aim of this survey is to provide insights to researchers into working of prefetching techniques and spark interesting future work for improving the performance advantages of prefetching even further.

A Survey of Techniques for Architecting and Managing GPU Register File

Sparsh Mittal, IEEE Transactions on Parallel and Distributed Systems (TPDS), 2016. http://dx.doi.org/10.1109/TPDS.2016.2546249

Abstract: To support their massively-multithreaded architecture, GPUs use very large register file (RF) which has a capacity higher than even L1 and L2 caches. In total contrast, traditional CPUs use tiny RF and much larger caches to optimize latency. Due to these differences, along with the crucial impact of RF in determining GPU performance, novel and intelligent techniques are required for managing GPU RF. In this paper, we survey the techniques for designing and managing GPU RF. We discuss techniques related to performance, energy and reliability aspects of RF. To emphasize the similarities and differences between the techniques, we classify them along several parameters. The aim of this paper is to synthesize the state-of-art developments in RF management and also stimulate further research in this area.

Algorithm-Directed Data Placement in Explicitly Managed Non-volatile Memory

P. Wu, D. Li, Z. Chen, Jeffrey Vetter, Sparsh Mittal, ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC), 2016. http://dx.doi.org/10.1145/2907294.2907321

Abstract: The emergence of many non-volatile memory (NVM) techniques is poised to revolutionize main memory systems because of the relatively high capacity and low lifetime power consumption of NVM. However, to avoid the typical limitation of NVM as the main memory, NVM is usually combined with DRAM to form a hybrid NVM/

Publications of Note (Continued)

DRAM system to gain the benefits of each. However, this integrated memory system raises a question on how to manage data placement and movement across NVM and DRAM, which is critical for maximizing the benefits of this integration. The existing solutions have several limitations, which obstruct adoption of these solutions in the high performance computing (HPC) domain. In particular, they cannot take advantage of application semantics, thus losing critical optimization opportunities and demanding extensive hardware extensions; they implement persistent semantics for resilience purpose while suffering large performance and energy overhead. In this paper, we re-examine the current hybrid memory designs from the HPC perspective, and aim to leverage the knowledge of numerical algorithms to direct data placement. With explicit algorithm management and limited hardware support, we optimize data movement between NVM and DRAM, improve data locality, and implement a relaxed memory persistency scheme in NVM. Our work demonstrates significant benefits of integrating algorithm knowledge into the hybrid memory design to achieve multi-dimensional optimization (performance, energy, and resilience) in HPC.

Communication Characterization and Optimization of Applications Using Topology-Aware Task Mapping on Large Supercomputers

Sarat Sreepathi, Ed D'Azevedo, Bobby Philip, and Patrick Worley. In Proceedings of the 7th ACM/SPEC on International Conference on Performance Engineering (ICPE '16). ACM, New York, NY, USA, 225-236, 2016. DOI=http://dx.doi.org/10.1145/2851553.2851575

Abstract: On large supercomputers, the job scheduling systems may assign a non-contiguous node allocation for user applications depending on available resources. With parallel applications using MPI (Message Passing Interface), the default process ordering does not take into account the actual physical node layout available to the application. This contributes to non-locality in terms of physical network topology and impacts communication performance of the application. In order to mitigate such performance penalties, this work describes techniques to identify suitable task mapping that takes the layout of the allocated nodes as well as the application's communication behavior into account. During the first phase of this research, we instrumented and collected performance data to characterize communication behavior of critical US DOE (United States - Department of Energy) applications using an augmented version of the mpiP tool. Subsequently, we developed several reordering methods (spectral bisection, neighbor join tree etc.) to combine node layout and application communication data for optimized task placement. We developed a tool called mpiAproxy to facilitate detailed evaluation of the various reordering algorithms without requiring full application executions. This work presents a comprehensive performance evaluation (14,000 experiments) of the various task mapping techniques in lowering communication costs on Titan, the leadership class supercomputer at Oak Ridge National Laboratory.

Anisotropic Electron-Photon and Electron-Phonon Interactions in Black Phosphorus

Xi Ling, Shengxi Huang, Eddwi H. Hasdeo, Liangbo Liang, William M. Parkin, Yuki Tatsumi, Ahmad R. T. Nugraha, Alexander A. Puretzky, Paul Masih Das, Bobby G. Sumpter, David Geohegan, Jing Kong, Riichiro Saito, Marija Drndic, Vincent Meunier, Mildred S. Dresselhaus, Nano Lett., 16,2260-2267 (2016). DOI: 10.1021/acs.nanolett.5b04540

Abstract: Orthorhombic black phosphorus (BP) and other layered materials, such as gallium telluride (GaTe) and tin selenide (SnSe), stand out among two-dimensional (2D) materials owing to their anisotropic in-plane structure. This anisotropy adds a new dimension to the properties of 2D materials and stimulates the development of angle-resolved photonics and electronics. However, understanding the effect of anisotropy has remained unsatisfactory to date, as shown by a number of inconsistencies in the recent literature. We use angle-resolved absorption and Raman spectroscopies to investigate the role of anisotropy on the electron–photon and electron–phonon interactions in BP. We highlight, both experimentally and theoretically, a nontrivial dependence between anisotropy and flake thickness and photon and phonon energies. We show that once understood, the anisotropic optical absorption appears to be a reliable and simple way to identify the crystalline orientation of BP, which cannot be determined from Raman spectroscopy without the explicit consideration of excitation wavelength and flake thickness, as commonly used previously.

Tailoring Plastics to Soak up Carbon Dioxide

Balaka Barkakaty, Katie L. Browning, Bobby Sumpter, David Uhrig, Ivana Karpisova, Kevin W. Harman, Ilia Ivanov, Dale K. Hensley, Jamie M. Messman, S. Michael Kilbey II, Bradley S. Lokitz; Macromolecules (2016) DOI: 10.1021/acs.macromol.5b02483.

Abstract: Development of novel polymeric materials capable of efficient CO2 capture and separation under ambient conditions is crucial for cost-effective and practical industrial applications. Here we report the facile synthesis of a new CO2-responsive polymer through postpolymerization modification of poly(2-vinyl-4,4-dimethylazlactone) (PVDMA). The reactive pendant azlactone groups of PVDMA are easily modified with 4-(N-methyltetrahydropyrimidine)benzyl alcohol (PBA) without any byproduct formation. FTIR and TGA experiments show the new PBA-functionalized polymer powder can reversibly capture CO2 at room temperature and under atmospheric pressure. CO2capture was selective, showing a high fixing efficiency even with a

Publications of Note (Continued)

mixed gas system (20% CO2, 80% N2) similar to flue gas. CO2 release occurred at room temperature, and release profiles were investigated as a function of temperature. Density functional theory (DFT) calculations coupled with modeling and simulation reveal the presence of two CO2 binding sites in the PBA-functionalized polymer resulting in a two-step CO2 release at room temperature. The ease of material preparation, high fixing efficiency, and robust release characteristics suggest that postpolymerization modification may be a useful route to designing new materials for CO2 capture.

Twisted MoSe2 Bilayers with Variable Local Stacking and Interlayer Coupling Revealed by Low-Frequency Raman Spectroscopy

Alexander A. Puretzky, Liangbo Liang, Xufan Li, Kai Xiao, Bobby G. Sumpter, Vincent Meunier, and David B. Geohegan; ACS Nano (2015). DOI:10.1021/acsnano.5b07807

Abstract: Unique twisted bilayers of MoSe2 with multiple stacking orientations and interlayer couplings in the narrow range of twist angles, 60 ± 3°, are revealed by low-frequency Raman spectroscopy and theoretical analysis. The slight deviation from 60° allows the concomitant presence of patches featuring all three high-symmetry stacking configurations (2H or AA′, AB′, and A′B) in one unique bilayer system. In this case, the periodic arrangement of the patches and their size strongly depend on the twist angle. Ab initio modeling predicts significant changes in frequencies and intensities of low-frequency modes versus stacking and twist angle. Experimentally, the variable stacking and coupling across the interface are revealed by the appearance of two breathing modes, corresponding to the mixture of the high-symmetry stacking configurations and unaligned regions of monolayers. Only one breathing mode is observed outside the narrow range of twist angles. This indicates a stacking transition to unaligned monolayers with mismatched atom registry without the in-plane restoring force required to generate a shear mode. The variable interlayer coupling and spacing in transition metal dichalcogenide bilayers revealed in this study may provide an interesting platform for optoelectronic applications of these materials.

Supramolecular polymerization of a prebiotic nucleoside provides insights into the creation of sequence-controlled polymers

Jun Wang, Peter V. Bonnesen, E. Rangel, E. Vallejo, Ariadna Sanchez-Castillo, H. James Cleaves II, Arthur P. Baddorf, Bobby G. Sumpter, Minghu Pan, Petro Maksymovych, and Miguel Fuentes-Cabrera; Scientific Reports (2015) DOI: 10.1038/srep18891

Abstract: A molecule, called a nucleoside analog and which is composed of an Adenine moeity and glycol group, was deposited on top of the Au(111) surface and studied with scanning tunneling microscopy and density functional theory calculations. The molecule was found to self-assemble into a supramolecular polymer, i.e., a polymer held by hydrogen bonds, that is not only two order of magnitudes longer than any other polymeric structure made by similar molecules on substrates but also encodes a sequence which is erasable with electron-induced excitation. The sequence encoded was the following: glycol-glycol-Adenine-Adenine-glycol-glycol, etc., and it remained intact for hundreds of nanometers. This work suggest that nucleoside polymerization on patterned substrates could be used to create sequence-controlled polymers that store information that is readable and erasable with microscopy techniques. Further, because polymerization takes place on a susbtrate, theory can be used to design what patterns and what molecules will produce the desired polymeric outcomes.

Monomer volume fraction profiles in pH responsive planar polyelectrolyte brushes

J. P. Mahalik, Y. Yang, Chaitra Deodhar, John F. Anker, Bradley S. Lokitz, S. Michael Kilbey, Bobby G. Sumpter, R. Kumar; J. Polym. Sci., Part B: Polym. Phys. 2016, 54, 956–964

Abstract: Spatial dependencies of monomer volume fraction profiles of pH responsive polyelectrolyte brushes were investigated using field theories and neutron reflectivity experiments. In particular, planar polyelectrolyte brushes in good solvent were studied and direct comparisons between predictions of the theories and experimental measurements are presented. The comparisons between the theories and the experimental data reveal that solvent entropy and ion-pairs resulting from adsorption of counterions from the added salt play key roles in affecting the monomer distribution and must be taken into account in modeling polyelectrolyte brushes. Furthermore, the utility of this physics-based approach based on these theories for the prediction and interpretation of neutron reflectivity profiles in the context of pH responsive planar polyelectrolyte brushes such as polybasic poly(2-(dimethylamino)ethyl methacrylate) (PDMAEMA) and polyacidic poly(methacrylic acid) (PMAA) brushes is demonstrated. The approach provides a quantitative way of estimating molecular weights of the polymers polymerized using surface-initiated atom transfer radical polymerization.

Stochastic finite element methods for partial differential equations with random input data

M. Gunzburger, C. Webster, and G. Zhang, "Stochastic finite element methods for partial differential equations with random input data", Acta Numerica, 2014 vol. 23 pp. 521-650.

Abstract: The quantification of probabilistic uncertainties in the outputs of physical, biological, and social systems

New Staff [Continued]

governed by partial differential equations with random inputs require, in practice, the discretization of those equations. Stochastic finite element methods refer to an extensive class of algorithms for the approximate solution of partial differential equations having random input data, for which spatial discretization is effected by a finite element method. Fully discrete approximations require further discretization with respect to solution dependences on the random variables. For this purpose several approaches have been developed, including intrusive approaches such as stochastic Galerkin methods, for which the physical and probabilistic degrees of freedom are coupled, and non-intrusive approaches such as stochastic sampling and interpolatory-type stochastic collocation methods, for which the physical and probabilistic degrees of freedom are uncoupled. All these method classes are surveyed in this article, including some novel recent developments. Details about the construction of the various algorithms and about theoretical error estimates and complexity analyses of the algorithms are provided. Throughout, numerical examples are used to illustrate the theoretical results and to provide further insights into the methodologies.

A Hyperspherical Adaptive Sparse-Grid Method for High-Dimensional Discontinuity Detection

G. Zhang, C. Webster, M. Gunzburger, and J. Burkardt, "A Hyperspherical Adaptive Sparse-Grid Method for High-Dimensional Discontinuity Detection," SIAM Journal on Numerical Analysis vol. 53 (3) pp. 1508-1536. DOI:10.1137/140971531

This work proposes and analyzes a hyperspherical adaptive hierarchical sparse-grid method for detecting jump discontinuities of functions in high-dimensional spaces. The method is motivated by the theoretical and computational inefficiencies of well-known adaptive sparse-grid methods for discontinuity detection. Our novel approach constructs a function representation of the discontinuity hypersurface of an N-dimensional discontinuous quantity of interest, by virtue of a hyperspherical transformation. Then, a sparse-grid approximation of the transformed function is built in the hyperspherical coordinate system, whose value at each point is estimated by solving a one-dimensional discontinuity detection problem. Due to the smoothness of the hypersurface, the new technique can identify jump discontinuities with significantly reduced computational cost, compared to existing methods. Moreover, hierarchical acceleration techniques are also incorporated to further reduce the overall complexity. Rigorous complexity analyses of the new method are provided as are several numerical examples that illustrate the effectiveness of the approach.

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