ISSUE
22
FALL
2016

QUARTERLY NEWSLETTER FOR THE COMPUTER SCIENCE AND MATHEMATICS DIVISION

CSMD

this issue

Science Highlights P.01
News P.15
Community Service P.15
Events P.16
Awards and Recognition P.17
Publications of Note P.17
About CSMD P.19

Science Highlights

A Cooperative Approach to Virtual Machine Based Fault Injection

Thomas Naughton, Christian Engelmann, Geoffroy Vallée, Ferrol Aderholdt, and Stephen L. Scott

Researchers developed a virtual machine (VM) based fault injection (FI) tool, which uses VM introspection techniques to overcome the "semantic gap" between the target (in VM) and FI controller (outside VM).

Resilience investigations often employ fault injection (FI) tools to study the effects of simulated errors on a target system. It is important to keep the target system under test (SUT) isolated from the controlling environment in order to maintain control of the experiment. Virtual machines (VMs) have been used to aid these investigations due to the strong isolation properties of system-level virtualization. A key challenge in fault injection tools is to gain proper insight and context about the SUT. In VM-based FI tools, this challenge of target context is increased due to the separation between host and guest (VM). We discuss an approach to VM-based FI that leverages virtual machine introspection (VMI) methods to gain insight into the target's context running within the VM. The key to this environment is the ability to provide basic information to the FI system that can be used to create a map of the target environment. We describe a proof-of-concept implementation and a demonstration of its use to introduce simulated soft errors into an iterative solver benchmark running in user-space of a guest VM.

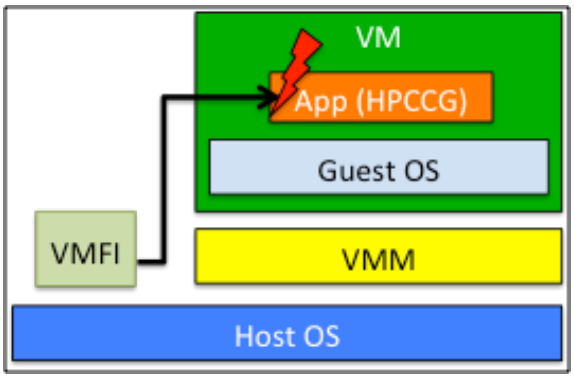

Illustration of the VM based fault-injection environment, with the target application (HPCCG) running in the Guest VM, and the FI utility running in the host using VM Introspection techniques to acquire the necessary context for the isolated target (in VM).

This work was supported by DOE/ASCR through the Hobbes OS/R project and previous exascale Virtualization projects.

Publication: T. Naughton, C. Engelmann, G. Vallée, F. Aderholdt, S.L. Scott, "A Cooperative Approach to Virtual Machine based Fault Injection", (To Appear) In Proceedings of the 22th European Conference on Parallel and Distributed Computing (Euro-Par) Workshops: 9th Workshop on Resiliency in High Performance Computing (Resilience) in Clusters, Clouds, and Grids Grenoble, France (Aug 23, 2016). Springer Verlag, Berlin, Germany.

OAK RIDGE National Laboratory

Science Highlights (Continued)

High-resolution Ensemble Projections of Near-term Regional Climate over the Continental United States

M. Ashfaq, M. D. Rastogi, R. Mei, S-C. Kao, S. Gangrade, B. Naz, D. Touma

In terms of the number of downscaled GCMs, horizontal grid spacing, and the length of simulations, the regional climate modeling in this study is one of the largest dynamical downscaling efforts over the continental U.S. to date.

We present high-resolution near-term ensemble projections of hydro-climatic changes over the contiguous U.S. using a regional climate model (RegCM4) that dynamically downscales 11 Global Climate Models from the 5th phase of Coupled Model Inter-comparison Project at 18km horizontal grid spacing. All model integrations span 41 years in the historical period (1965–2005) and 41 years in the near-term future period (2010–2050) under Representative Concentration Pathway 8.5 and cover a domain that includes the contiguous U.S. and parts of Canada and Mexico. Should emissions continue to rise, surface temperatures in every region within the U.S. will reach a new climate norm well before mid 21st century regardless of the magnitudes of regional warming. Significant warming will likely intensify the regional hydrological cycle through the acceleration of the historical trends in cold, warm and wet extremes. The future temperature response will be partly regulated by changes in snow hydrology over the regions that historically receive a major portion of cold season precipitation in the form of snow. Our results indicate the existence of the Clausius-Clapeyron scaling at regional scales where per degree centigrade rise in surface temperature will lead to a 7.4% increase in precipitation from extremes. More importantly, both winter (snow) and summer (liquid) extremes are projected to increase across the U.S. These changes in precipitation characteristics will be driven by a shift towards shorter and wetter seasons. Overall, projected changes in the regional hydro-climate can have substantial impacts on the natural and human systems across the U.S.

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

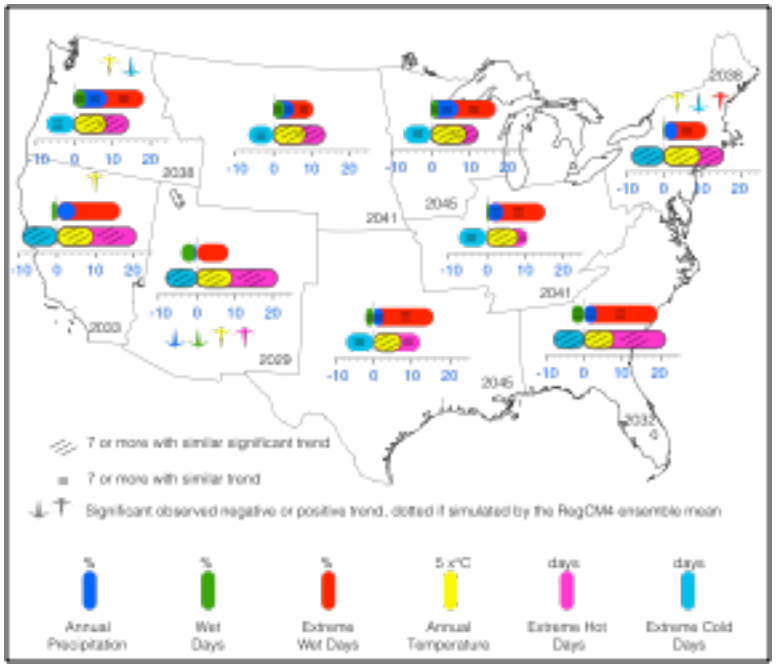

Summary of observed and near-term projected mean and extreme precipitation and temperature trends over the U.S. climate regions

Publication: M. Ashfaq, M. D. Rastogi, R. Mei, S-C. Kao, S. Gangrade, B. Naz, D. Touma (2016), High-resolution ensemble projections of near-term regional climate over the continental United States, JGR-Atmospheres, accepted.

Science Highlights (Continued)

Regional Hydrologic Response to Climate Change over the Conterminous United States

B.S. Naz, S-C. Kao, M. Ashfaq, D. Rastogi, R. Mei, L. C. Bowling

Using a hierarchical modeling framework, this research provides one of the most detailed (to date) high-resolution (4km) resolution hydrological projections over the conterminous United States for the near-term future.

Despite the fact that Global Climate Model (GCM) outputs have been used to project hydrologic impacts of climate change using off-line hydrologic models for two decades, many of these efforts have been disjointed — applications or at least calibrations have been focused on individual river basins and using a few of the available GCMs. This study improves upon earlier attempts by systematically projecting hydrologic impacts for the entire conterminous United States (US), using outputs from ten GCMs from the latest Coupled Model Intercomparison Project phase 5 (CMIP5) archive, with seamless hydrologic model calibration and validation techniques to produce a spatially and temporally consistent set of current hydrologic projections. The Variable Infiltration Capacity (VIC) model was forced with ten-member ensemble projections of precipitation and air temperature that were dynamically downscaled using a regional climate model (RegCM4) and bias-corrected to 1/24° (~ 4 km) grid resolution for the baseline (1966–2005) and future (2011–2050) periods under the Representative Concentration Pathway 8.5. Based on regional analysis, the VIC model projections indicate an increase in winter and spring total runoff due to increases in winter precipitation of up to 20% in most regions of the US. However, decreases in snow water equivalent (SWE) and snow-covered days will lead to significant decreases in summer runoff with more pronounced shifts in the time of occurrence of annual peak runoff projected over the eastern and western US. In contrast, the central US will experience year-round increases in total runoff, mostly associated with increases in both extreme high and low runoff. The projected hydrological changes described in this study have implications for various aspects of future water resource management, including water supply, flood and drought preparation, and reservoir operation.

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

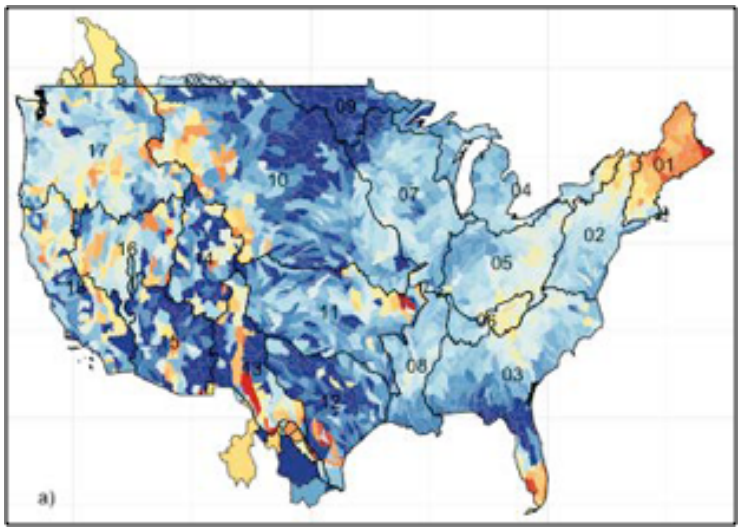

Ensemble median percentage changes in high runoff (the 95th percentile of daily runoff) in 2011-2050 w.r.t. 1966-2005. The numbers refer to the HUC2 identification of hydrologic regions.

Publication: B.S. Naz, S-C. Kao, M. Ashfaq, D. Rastogi, R. Mei, L. C. Bowling (2016), Regional hydrologic response to climate change in the conterminous United States using high-resolution hydroclimate simulations, doi:10.1016/j.gloplacha.2016.06.003

Science Highlights (Continued)

Petascale Simulations Expose a Rational Tunability of Bulk Heterostructure Morphology and the Molecular Interfaces

Jan-Michael Y. Carrillo, Zach Seibers, Rajeev Kumar, Michael A. Matheson, John F. Ankner, Monojoy Goswami, Kiran Bhaskaran-Nair, William A. Shelton, Bobby G. Sumpter, and S. Michael KilbeyII

Petascale molecular dynamics, atomistic, and density functional computation reveal the evolution of internal mixing of poly(3-hexylthiophene) (P3HT) and phenyl-C61-butyric acid methyl ester (PCBM) allowing comparison with neutron-reflectivity-derived density profiles. Controlling nanoscale mixing of organic photovoltaic (OPV) donor and acceptor materials is key to optimizing their performance. This work serves as a roadmap for rational design of conjugated-polymer based composite materials.

Understanding how additives interact and segregate within bulk heterojunction (BHJ) thin films is critical for exercising control over structure at multiple length scales and delivering improvements in photovoltaic performance. The morphological evolution of poly(3-hexylthiophene) (P3HT) and phenyl-C61-butyric acid methyl ester (PCBM) blends that are commensurate with the size of a BHJ thin film was examined using petascale computational simulations. Comparisons between two-component and three-component systems containing short P3HT chains as additives undergoing thermal annealing demonstrated that short chains alter the morphology in useful ways: they efficiently migrate to the P3HT/PCBM interface, increasing the P3HT domain size and interfacial area. Simulation results agree with depth profiles determined from neutron reflectometry measurements that reveal PCBM enrichment near substrate and air interfaces but a decrease in that PCBM enrichment when a small amount of short P3HT chains are integrated into the BHJ blend. Atomistic simulations of the P3HT/PCBM blend interfaces show a non-monotonic dependence of the interfacial thickness as a function of number of repeat units in the oligomeric P3HT additive, and the thiophene rings orient parallel to the interfacial plane as they approach the PCBM domain. Using the nanoscale geometries of the P3HT oligomers, LUMO and HOMO energy levels calculated by density functional theory are found to be invariant across the donor/acceptor interface. These connections between additives, processing, and morphology at all length scales are generally useful for efforts to improve device performance.

We thank C. Rycroft for assistance with radical Voronoi analysis using MPI and Voro++, A.V. Dobrynin for assistance with the development of the substrate–BHJ interactions, T.D. Nguyen and W.M. Brown for assistance with LAMMPS MD code, and M. Chen for assistance with NWChem code. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory (ORNL), which is supported by the Office of Science of the U.S. Department of Energy (DoE) under Contract No. DEAC05-00OR22725. This research was

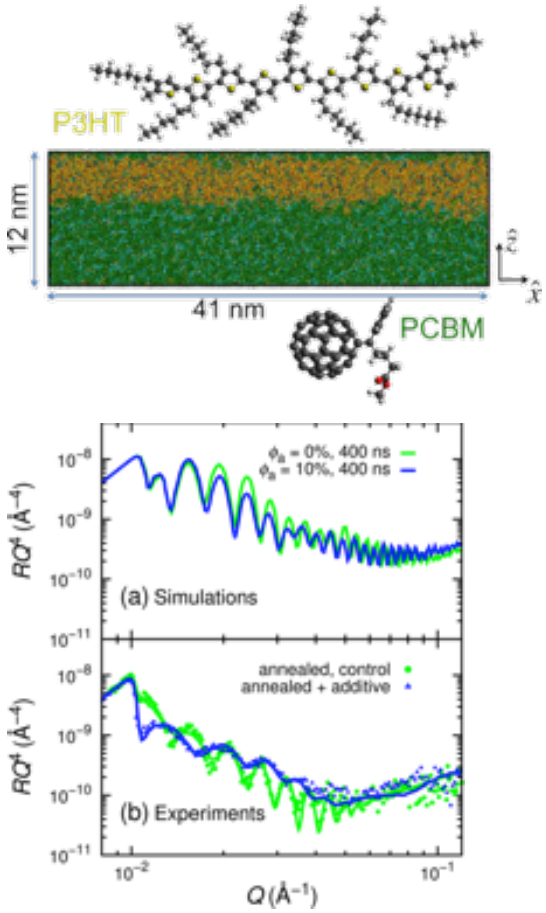

Coarse-grained simulation of OPV interface (top) compared with neutron reflectivity measurements (bottom).

conducted in part at the Center for Nanophase Materials Sciences (CNMS) and at the Spallation Neutron Source (SNS), which are sponsored at ORNL by the Scientific User Facilities Division, Office of Basic Energy Sciences (BES), U.S. DoE. Experimental work of Z.S. and S.M.K. was sponsored by National Science Foundation under Award No. EPS 1004083 and Award No. 1512221. M.G. and J.-M.Y.C. acknowledge support by the U.S. DoE, BES, Materials Science and Engineering Division (MSED). The computational work conducted by K.B.-N. and W.A.S. is supported by the U.S. DoE under EPSCoR Grant No. DESC0012432 with additional support from the Louisiana Board of Regents.

Publication: "Petascale Simulations of the Morphology and the Molecular Interface of Bulk Heterojunctions" Jan-Michael Y. Carrillo, Zach Seibers, Rajeev Kumar, Michael A. Matheson, John F. Ankner, Monojoy Goswami, Kiran Bhaskaran-Nair, William A. Shelton, Bobby G. Sumpter, and S. Michael KilbeyII; ACS Nano (2016) DOI: 10.1021/acsnano.6b03009

Science Highlights (Continued)

Havens: Explicit Reliable Memory Regions for HPC Applications

S. Hukerikar, and C. Engelmann

Researchers developed a memory management method based on the notion of memory regions that prioritizes resilience and locality for high-performance computing (HPC) applications. Defined an abstract interface and developed a library-based proof-of-concept prototype implementation. Demonstrated its viability with a conjugate gradient solver and memory regions with different resilience against memory errors. Created programming language annotations to make a program's Havens more explicit and to easily incorporate resilience capabilities in HPC application codes.

Havens provide robust memory regions in which program objects may be allocated. A memory region is protected by a predefined error detection and/or correction scheme that is agnostic to the algorithm features or structure of the data. This approach to memory management enables separation between the memory allocation and the implementation of the robustness scheme. Havens enable applications to exert fine-grained control on the resilience properties of individual program objects. Since different Havens may have varying guarantees of reliability, based on the strength of the protection mechanism and its performance overhead, object placement in Havens may be driven by the trade-off between criticality of the object to program correctness and the associated overhead.

The language support for Havens offers explicit, convenient, sound and scalable control over Havens. The typing system enables application programmers to statically encode memory management decisions on the basis of their understanding of their application resilience requirements. This enables 3 different application-level reliability models: (1) selective reliability with Havens providing selective regions of program memory with comprehensive error protection, (2) specialized reliability with Havens providing different protection schemes based on application resilience and performance needs, and (3) dynamic reliability with Havens being enabled/disabled on demand during application execution.

Through accelerated fault injection experiments on a preconditioned conjugate gradient (PCG) solver (below (a)), we observed that the allocation of the static state into Havens yields the highest rate of successful completion. Also, providing the highest fault protection coverage does not yield a comparable increase in application resilience. Therefore, imposing the selective reliability model using Havens on the PCG solver provides the right balance between application resilience and performance overhead (below (b)).

This work was sponsored by the US Department of Energy, Office of Science, Early Career Research Program.

Publications:

- S. Hukerikar and C. Engelmann, "Havens: Explicit Reliable Memory Regions for HPC Applications", 20th IEEE High Performance Extreme Computing Conference (HPEC) 2016, Waltham, MA, September 13-15, 2016.
- S. Hukerikar and C. Engelmann, " Language Support for Reliable Memory Regions", 29th International Workshop on Languages and Compilers for Parallel Computing (LCPC) 2016, Rochester, NY, September, 28-30, 2016.

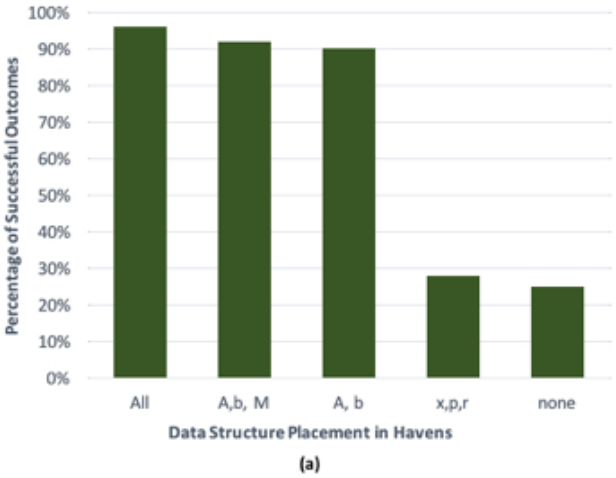

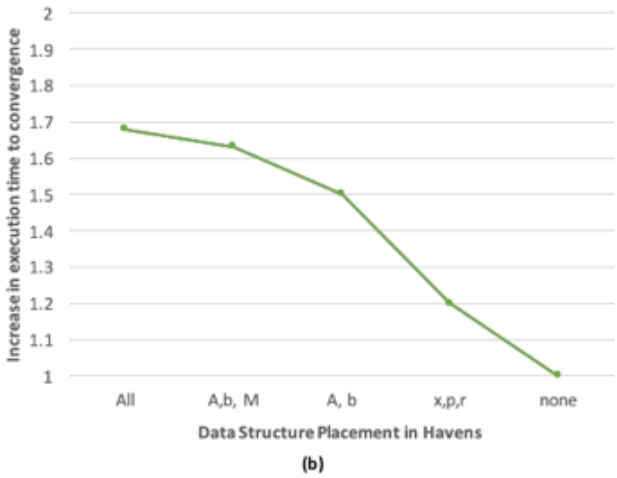

Fault injection campaign results using Havens with a preconditioned conjugate gradient solver

Science Highlights (Continued)

Improving the Functionality of Atomically Thin Layered Materials

Jian Gao, Young Duck Kim, Liangbo Liang, Juan Carlos Idrobo, Phil Chow, Jiawei Tan, Baichang Li, Lu Li, Bobby G. Sumpter, Toh-Ming Lu, Vincent Meunier, James Hone, Nikhil Koratkar

This work enables stable transition-metal substitution that preserves the monolayer's semiconducting nature, along with other characteristics, including direct bandgap photoluminescence. Such tuning of functionality achieves one major prerequisite for integration into modern solid-state electronic and optoelectronic technology.

As device length scales progressively shrink, they push the limits of silicon technology, and new challenges emerge that must be addressed to meet this requirement. The family of layered semiconductors, known as the transition metal dichalcogenides (TMDCs), represents one such alternative that has met a surge of recent interest for optoelectronics applications due to attractive physical scaling properties and exotic behavior in the atomically thin regime. However appropriate doping and control has been largely lacking. In this work, we successfully demonstrate large-area "in situ" transition-metal doping for synthetic TMDC monolayers achieved by doped atomically thin WS_2 and MoS_2 monolayers with niobium (Nb) and rhenium (Re) impurities during the CVD-growth process. This method is based on a thermal chemical vapor deposition (CVD) approach that entails a scalable one-pot synthesis procedure. Detailed characterization and theoretical modeling show that the situ doping of transition-metal atoms during the growth of chemical vapor deposited monolayer MoS_2 and WS_2, leave the direct-gap semiconductor behavior of the atomically thin materials intact. Thus our results showcase a powerful doping strategy for altering the electronic behavior of atomically thin, chemical vapor deposited TMDCs by substitutional doping of the transition-metal atoms with impurity transition metals.

N.K. acknowledges funding support from the USA National Science Foundation (Award Numbers: 1435783, 1510828 and 1608171), New York State under NYSTAR program C080117, and from the John A. Clark and Edward T. Crossan Endowed Chair Professorship at the Rensselaer Polytechnic Institute (RPI). The theoretical work was also supported in part by the Office of Naval Research. The computations were performed using the resources of the Center for Computational Innovation at RPI. L.L. was supported as a Eugene P. Wigner Fellow at Oak Ridge National Laboratory. J.H. and Y.D.K. acknowledge support from the NSF MRSEC program through Columbia in the Center for Precision Assembly of Superstratic and Superatomic Solids (DMR-1420634). Microscopy research was conducted as part of a user proposal through ORNL's Center for Nanophase Materials Sciences, which is a U.S. Department of Energy, Office of Science User Facility (J.C.I.).

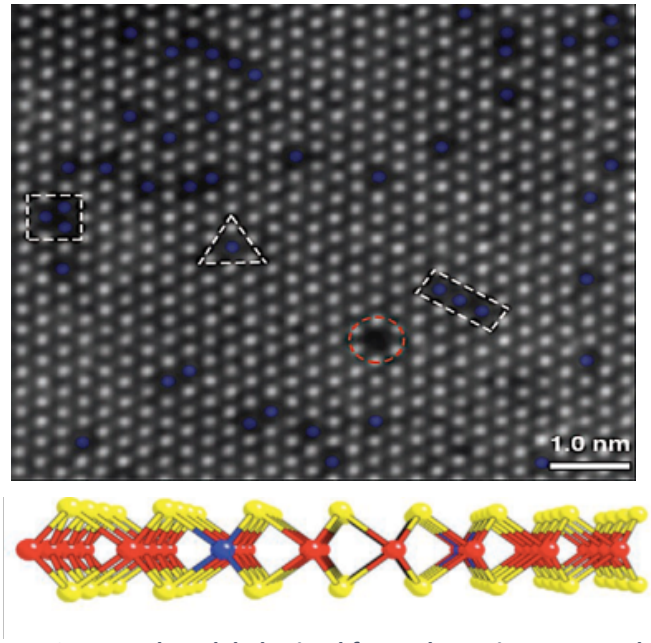

Top: Structural model obtained from Aberration-corrected Z-contrast STEM image of monolayer Nb-doped WS_2 showing the distribution of Nb dopants (blue, Nb sites; white, W sites; black hole, metal vacancy). The Nb concentration is estimated as ≈6.7%. Bottom: Cross-section schematic of monolayer WS_2 with Nb doping (W, red; Nb, blue; sulfur, yellow).

Publication: "Transition-Metal Substitutional Doping in Synthetic Atomically Thin Semiconductors" Jian Gao, Young Duck Kim, Liangbo Liang, Juan Carlos Idrobo, Phil Chow, Jiawei Tan, Baichang Li, Lu Li, Bobby G. Sumpter, Toh-Ming Lu, Vincent Meunier, James Hone, Nikhil Koratkar; Adv. Mater. (2016) DOI: 10.1002/adma.201601104

Science Highlights (Continued)

A Hybrid Computational Strategy to Address WGS Variant Analysis in >5000 Samples

Zhuoyi Huang, Navin Rustagi, Narayanan Veeraraghavan, Andrew Carroll, Richard A. Gibbs, Eric Boerwinkle, Manjunath Gorentla Venkata, Fuli Yu

Researchers developed scalable hybrid computational strategy to conduct analysis of large-scale whole genome datasets.

The decreasing costs of sequencing are driving the need for cost effective and real time variant calling of whole genome sequencing data. The scales of these projects are far beyond the capacity of typical computing resources available with most research labs. Other infrastructures like the cloud AWS environment and supercomputers also have limitations due to which large scale joint variant calling becomes infeasible, and infrastructure specific variant calling strategies either fail to scale up to large datasets or abandon joint calling strategies. In this paper, we present a high throughput framework including multiple variant callers for single nucleotide variant (SNV) calling, which leverages hybrid-computing infrastructure consisting of cloud AWS, supercomputers and local high performance computing infrastructures. We present a novel binning approach for large scale joint variant calling and imputation which can scale up to over 10,000 samples while producing SNV callsets with high sensitivity and specificity. As a proof of principle, we present results of analysis on Cohorts for Heart And Aging Research in Genomic Epidemiology (CHARGE) WGS freeze 3 dataset in which joint calling, imputation and phasing of over 5300 whole genome samples was produced in under 6 weeks using four state-of-the-art callers.

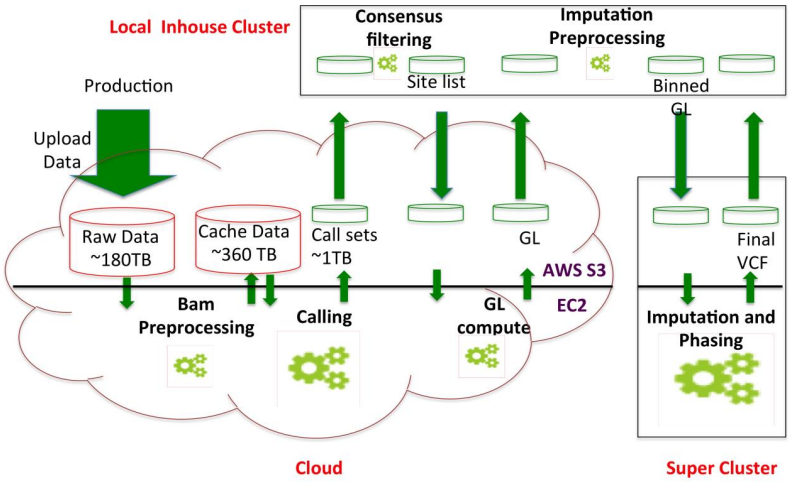

In the news:

HPCWire:

https://www.hpcwire.com/2016/09/22/new-genomics-pipeline-combines-aws-local-hpc-supercomputing/

Baylor College of Medicine News:

https://www.bcm.edu/news/genome-sequencing/new-scalable-whole-genome-data-analysis

ORNL News:

https://www.ornl.gov/news/ornl-helps-develop-hybrid-computational-strategy-efficient-sequencing-massive-genome-datasets

Gaining Control of Molecular Ionic Nanostructures

Seokmin Jeon, Peter Doak, Bobby G. Sumpter, Panchapakesan Ganesh, and Petro Maksymovych

New two-dimensional molecular ionic structures TTFx-TCNQy (x = 1..2, y=1..13) grow on noble metal surfaces. Multiplicity of stable molecular 2D structures, their rich phase diagram and the corresponding phase-transitions present a new opportunity for low-dimensional molecular systems with strong electron correlations.

Bulk molecular ionic solids exhibit fascinating electronic properties, including electron correlations, phase transitions and superconducting ground states. In contrast, few of these phenomena have so far been observed in low-dimensional molecular structures, including thin films, nanoparticles and molecular blends, not in the least because most of such structures have so far been composed of nearly closed-shell molecules. In this work we show the existence of the surface phase-diagram that controls the structures of tetrathiafulvalene (TTF) donor and 7,7,8,8-tetracyanoquinodimethane (TCNQ) acceptor molecules on the surfaces, and demonstrate phase-transitions that occur upon progressively increasing the density of TCNQ while keeping the surface coverage of TTF fixed. A binding motif that underlies the stable phases and infer the dominant interactions that enable the existence of the rich spectrum of surface structures is inferred.

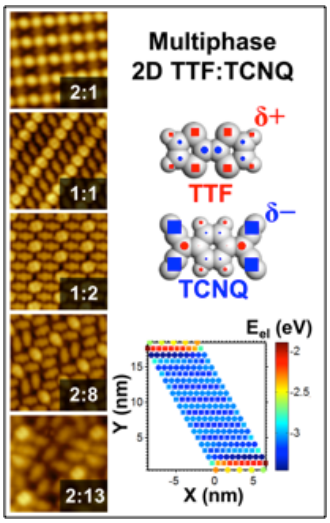

Multiphases exhibited by TTF and TCNQ on a silver surface.

This work was conducted at the Center for Nanophase Materials Sciences, which is a DOE Office of Science User Facility. S.J. thanks Geoffrey A. Rojas for sharing preliminary STM results and Changwon Park for helpful discussions about theoretical computations.

Publication: "Thermodynamic Control of Two-Dimensional Molecular Ionic Nanostructures on Metal Surfaces", Seokmin Jeon, Peter Doak, Bobby G. Sumpter, Panchapakesan Ganesh, Petro Maksymovych, Center for Nanophase Materials Sciences, ORNL, ACS Nano 10, 7821–7829 (2016). DOI: 10.1021/acsnano.6b03492

Science Highlights (Continued)

An Electrolyte-free Polymeric "smart material"

Jian Gao, Young Duck Kim, Liangbo Liang, Juan Carlos Idrobo, Phil Chow, Jiawei Tan, Baichang Li, Lu Li, Bobby G. Sumpter, Toh-Ming Lu, Vincent Meunier, James Hone, Nikhil Koratkar

Researchers produced an electrolyte-free polymeric "smart material" based on a thermo-controllable nanostructure.

This study shows promise for enabling more flexibility in processing by utilizing water-soluble conjugated polymers in aqueous solutions for self-assembly based fabrication of stimuli-responsive nanostructures and sensory materials. The structural and electrical properties of the conjugated polymer-embedded self-assembled architectures were investigated by combining small-angle neutron and X-ray scattering, coarse-grained molecular dynamics simulations, and impedance spectroscopy. Results show how a pluronic L62 solution template organizes the water soluble conjugated polymers by stably incorporating them into the hydrophilic domains which effectively inhibits the typical aggregation. The thermo-tunable morphology of L62 from the micellar-to-lamellar phase defines the embedded conjugated polymer network. As a result, the conductivity is strongly coupled to the structural change of the templating L62 phase and exhibits a thermally reversible behavior with no signs of quenching of the conductivity at high temperature.

The Research at Oak Ridge National Laboratory's Spallation Neutron Source and Center for Nanophase Materials Sciences was sponsored by the Scientific User Facilities Division, Office of Basic Energy Sciences, U.S. Department of Energy (DOE). Research was also sponsored by the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the U. S. Department of Energy. The authors thank Dr. Sungsik Lee for his assistance during the experiments conducted at 12-BM-B of the Advanced Photon Source, Argonne National Laboratory. Z.Z. gratefully acknowledges the financial support from Jülich Center for Neutron Science, Research center Jülich. This research used resources of the Oak Ridge Leadership Computing Facility (OLCF), which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725. J. -M. Y. C. acknowledges the Center for Accelerated Materials Modeling (CAMM) funded by the US DoE, BES, Materials Science and Technology Division (MSTD). Y.L. and M.P.P acknowledge the support for conductivity measurements provided by the US DOE, Office of Science, BES, Materials Science and Engineering Division.

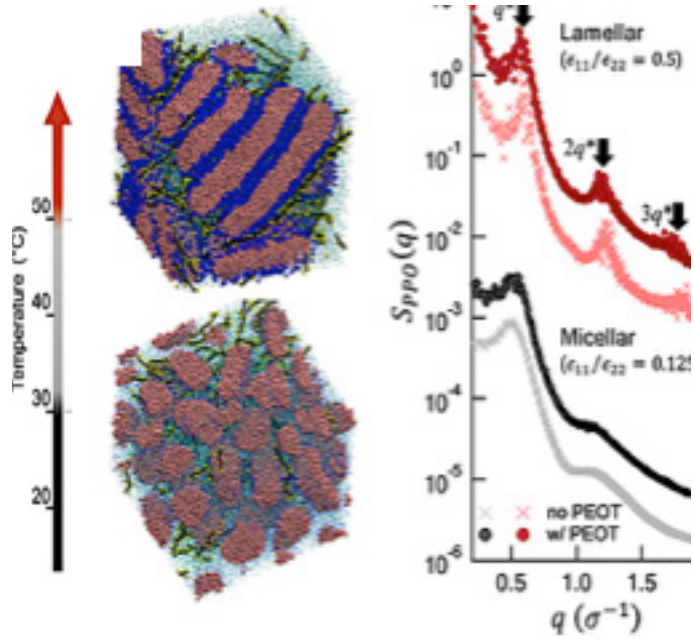

Thermotropic morphology changes. PEO (blue) and PPO (pink); D2O are the cyan. Lamellar peaks at q*, 2q*, and 3q* (q*≈0.6 σ-1) with increasing temperature are clearly observed.

Publication: Jian Gao, Young Duck Kim, Liangbo Liang, Juan Carlos Idrobo, Phil Chow, Jiawei Tan, Baichang Li, Lu Li, Bobby G. Sumpter, Toh-Ming Lu, Vincent Meunier, James Hone, Nikhil Koratkar, "Transition-Metal Substitutional Doping in Synthetic Atomically Thin Semiconductors", Small. (2016) DOI: 10.1002/smll.201601342

Science Highlights (Continued)

Multiphase MoS2 Monolayers as Efficient Photocatalysts for Hydrogen Production from Water Splitting

R. Peng, L. Liang, Z. D. Hood, A. Boulesbaa, A. Puretzky, A. V. Ievlev, J. Come, O. S. Ovchinnikova, H. Wang, C. Ma, M. Chi, B. G. Sumpter, Z. Wu

2D multiphasic MoS2 nanosheets were synthesized and tailored as visible light photocatalysts for efficient solar hydrogen production from water splitting. This development of efficient and inexpensive photocatalysts that don't involve noble metals is crucial for harvesting solar energy into fuels.

This work shows how the confluence of unique optical, electronic and chemical properties of 2D MoS2 monolayers can be fully captured for efficient photocatalytic water splitting. 2D single-layer MoS2 nanosheets are demonstrated as efficient photocatalysts for hydrogen evolution reaction (HER) from water splitting thanks to specific in-plane heterojunctions constructed in the MoS2 monolayer. These functional heterojunctions are formed among the different phases of chemically exfoliated MoS2 monolayers: semiconducting 2H, metallic 1T, and nearly-metallic 1T' phases. The proportion of the three MoS2 phases can be systematically controlled via thermal annealing of the nanosheets. Interestingly, a volcano relationship is observed between the photocatalytic HER activity and the annealing temperature with an optimum activity obtained after annealing at 60ºC. The presence of 2H/1T' heterostructures in the 60 oC treated MoS2 is responsible for the efficient separation of photogenerated charge carriers and abundant active basal sites for hydrogen evolution, as supported by results from Raman, XPS, conductive AFM, time-resolved spectroscopy, and the plane-wave DFT calculations. It clearly demonstrated that MoS2 nanosheets can be rationally engineered via phase manipulation to enable efficient hydrogen evolution from water splitting under visible light irradiation.

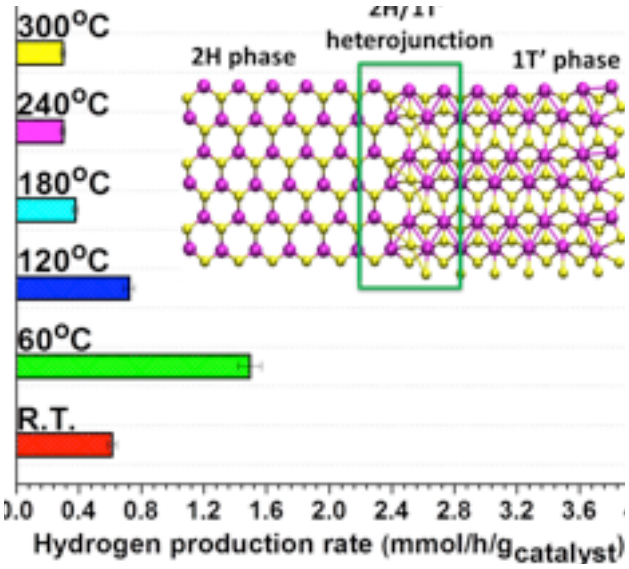

Multiphasic MoS2 treated at 60°C synergizes the photocatalyst (2H) and co-catalyst (1T') and enhances the photocatalytic solar hydrogen production.

This research was conducted at the Center for Nanophase Materials Sciences, which is a DOE Office of Science User Facility. LL was supported by Eugene P. Wigner Fellowship at ORNL. ZDH acknowledges support from the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1148903 and the Georgia Tech-ORNL Fellowship.

Publication: "In-plane heterojunctions enable multiphasic 2D MoS2 nanosheets as efficient photocatalysts for hydrogen evolution from water splitting", R. Peng, L. Liang, Z. D. Hood, A. Boulesbaa, A. Puretzky, A. V. Ievlev, J. Come, O. S. Ovchinnikova, H. Wang, C. Ma, M. Chi, B. G. Sumpter, Z. Wu

Science Highlights (Continued)

Evaluating OpenSHMEM Explicit Remote Memory Access Operations and Merged Requests

Swen Boehm, Swaroop Pophale and Manjunath Gorentla Venkata

Researchers implemented and evaluated merged handle approach for aggregating related RMA operations.

As we often see in scientific code, a series of updates are made during the computation phase and are written during the communication phase. Most updates need to happen together to enable the next set of computations. Such updates can be merged together to enable easy checking for the user. This has many performance as well as productivity implications. This approach may greatly simplify how users write their code, replacing multiple request handles by a single request handle. The performance advantage comes from the fact that testing completion of a single handle is much more cost efficient than using either checking individual handles or executing mass memory updates via quiet or barrier that will only return after all pending local and remote memory updates are processed.

We introduce non-blocking RMA operations with requests, where each operation has an explicit request to track and complete the operation. Second, we introduce interfaces to merge multiple requests into a single request handle. The merged request tracks multiple user-selected RMA operations, which provides the flexibility of tracking related communication operations with one request handle. Lastly, we explore the implications in terms of performance, productivity, usability and the possibility of defining different patterns of communication via merging of requests. Our experimental results show that a well-designed and implemented OpenSHMEM stack can hide the overhead of allocating and managing the requests. The latency of RMA operations with requests is similar to blocking and implicit non-blocking RMA operations. We test our implementation with the Scalable Synthetic Compact Applications (SSCA #1) benchmark and observe that using RMA operations with requests and merging of these requests outperform the implementation using blocking RMA operations and implicit non-blocking operations by 49% and 74% respectively.

Work was supported by the DoD and the Extreme Scale Systems Center at ORNL

Publication: Swen Boehm, Swaroop Pophale and Manjunath Gorentla Venkata, “Evaluating OpenSHMEM Explicit Remote Memory Access Operations and Merged Requests” In Proceedings of Third OpenSHMEM Workshop, Baltimore, MD, USA, August 2016.

A Computational Workflow for Silicon Donor Qubits

Travis S. Humble, M. Nance Ericson, Jacek Jakowski, Jingsong Huang, Charles Britton, Franklin G. Curtis, Eugene F. Dumitrescu, Fahd A. Mohiyaddin, and Bobby G. Sumpter

Researchers demonstrated a unique computational workflow for characterizing the behavior of qubit developed from phosphorous-doped silicon devices. This work provides a first-in-kind tool for evaluating the influence of material properties on the quantum computing devices built using CMOS-compatible technology.

Developing devices that can reliably and accurately demonstrate the principles of superposition and entanglement is an on-going challenge for the quantum computing community. Modeling and simulation offer attractive means of testing early device designs and establishing expectations for operational performance. However, the complex integrated material systems required by quantum device designs are not captured by any single existing computational modeling method. We examine the development and analysis of a multi-staged computational workflow that can be used to design and characterize silicon donor qubit systems with modeling and simulation. Our approach integrates quantum chemistry calculations with electrostatic field solvers to perform detailed simulations of a phosphorus dopant in silicon. We show how atomistic details can be synthesized into an operational model for the logical gates that define quantum computation in this particular technology. The resulting computational workflow realizes a design tool for silicon donor qubits that can help verify and validate current and near-term experimental devices.

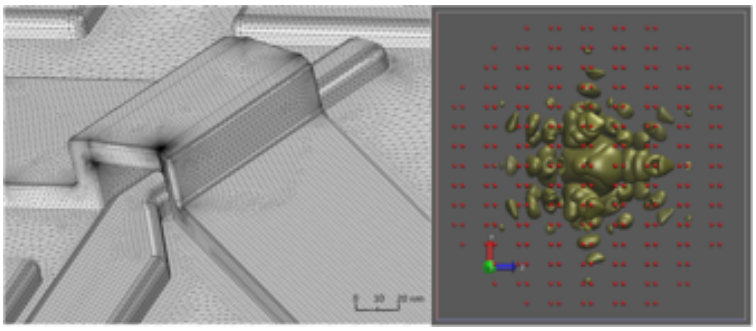

(Left) Nanoscale electronic circuit to control a single phosphorous atom; (Right) electron spin density for phosphorous-doped nanocrystal

This work was sponsored by the ORNL Laboratory Directed Research and was supported in part by the CNMS.

Publication: T. S. Humble, M. N. Ericson, J. Jakowski, J. Huang, C. Britton, F. G. Curtis, E. Dumitrescu, F. A. Mohiyaddin, and B. G. Sumpter, “A Computational Workflow for Designing Silicon Donor Qubits,” Nanotechnology 27, 424002 (2016); cf. arXiv:1608.01416

Science Highlights (Continued)

OpenSHMEM Implementation of HPCG Benchmark

E. D'Azevedo, S. Powers, and N. Imam

Researchers ported the High Performance Conjugate Gradient (HPCG) benchmark to an OpenSHMEM and MPI one-sided implementation. This work demonstrates that an OpenSHMEM implementation of HPCG is able to obtain similar performance to the original MPI version.

We describe the effort to implement the HPCG benchmark using OpenSHMEM and MPI one-sided communication. Unlike the High Performance LINPACK (HPL) benchmark that places emphasis on large dense matrix computations, the HPCG benchmark is dominated by sparse operations such as sparse matrix-vector product, sparse matrix triangular solve, and long vector operations. The MPI one-sided implementation is developed using the one-sided OpenSHMEM implementation. Preliminary results comparing the original MPI, OpenSHMEM and MPI one-sided implementations on an SGI cluster, Cray XK7 and Cray XC30 are presented. The results suggest the MPI, OpenSHMEM and MPI one-sided implementations all obtain similar overall performance but the MPI one-sided implementation seems to slightly increase the run time for multigrid preconditioning in HPCG on the Cray XK7 and Cray XC30.

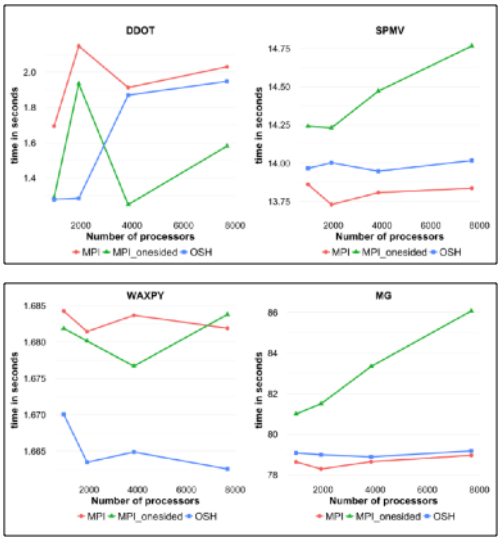

Performance of HPCG using 32 MPI tasks on each node (Cray XC30)

This work is supported by the DoD, the Extreme Scale Systems Center at ORNL, and the Oak Ridge Leadership Computing Facility

Publication: E. D'Azevedo, S. Powers and N. Imam, "OpenSHMEM Implementation of HPCG Benchmark," in Proceedings of Third OpenSHMEM Workshop, Baltimore, MD, USA, August 2016.

Designing Functional Molecules to Remove Interactions that Cause Aggregation

Hong-Hai Zhang, Chuanxu Ma, Peter V. Bonnesen, Jiahua Zhu, Bobby G. Sumpter, Jan-Michael Y. Carrillo, Panchao Yin, Yangyang Wang, An-Ping Li, and Kunlun Hong

For the first time, researchers were able to synthesize Poly(5-alkyl-2,3-thiophene)s (P5AT), an ortho-connected isomer of poly(3-alkyl-2,5-thiophene)s (P3AT). P5ATs show distinct helical structures with different optoelectronic properties compared to their P3AT isomeric counterparts and are representative of a new class of potential optoelectronic materials.

Conjugated polyarylenes have found a wide range of applications including optical sensors, electrochromic devices, field effect transistors, and organic solar cells. However often the interactions driven by the conjugated pi-backbone causes difficulties in solution processing due to strong aggregation of the polymer chains. One approach to uncouple these backbone interactions between polymer chains to mediate aggregation is to introduce a helix geometry for the backbone. This could be done by using ortho-connected isomers, poly(5-alkyl-2,3-thiophene)s (P5AT) but such polyarylenes have never been reported because of the difficulty in their syntheses. This work marks the first synthesis of regioregular P5AT; achieved via controlled Suzuki cross-coupling polymerization with PEPPSI-IPr as catalyst, affording the polymers with tunable molecular weight, narrow polydispersity (PDI), and well-defined functional end groups at the gram scale. The helical geometry of the P5ATs predicted by computational modeling was confirmed by a combination of NMR, small-angle X-ray scattering (SAXS), and scanning tunneling microscopy (STM). Since the Suzuki cross-coupling polymerization tolerates a wide range of functionalized aromatic monomers, this method can be applicable for

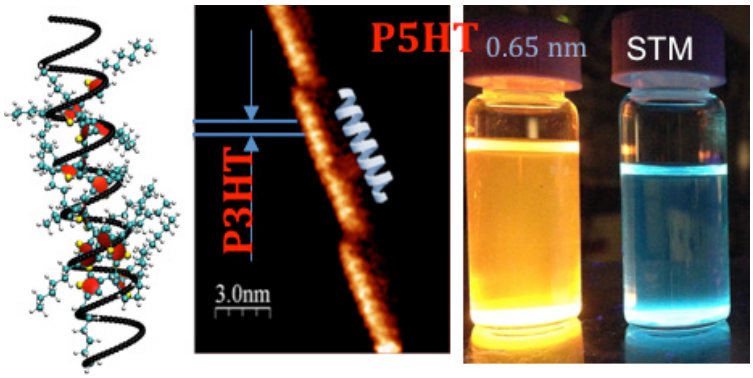

Left: Computational modeling predicted structure for P5BT. Middle: STM image of P5BT on HOPG showing a 0.65 nanometer helical pitch. Right: Photographs of the THF solutions of P3HT and P5HT under 360 nm UV light showing that P5HT has green photoluminescence under irradiation.

the synthesis of poly(o-arylene)s with aromatic rings other than thiophene and thus opens access to a new class of poyarylenes where the combination of helical structure, conjugated backbone, and photoluminescence properties make P5ATs good candidates for a range of applications in optoelectronics and related field applications.

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

Publication: "Helical Poly(5-alkyl-2,3-thiophene)s: Controlled Synthesis and Structure Characterization"; Hong-Hai Zhang, Chuanxu Ma, Peter V. Bonnesen, Jiahua Zhu, Bobby G. Sumpter, Jan-Michael Y. Carrillo, Panchao Yin, Yangyang Wang, An-Ping Li, and Kunlun Hong

Science Highlights (Continued)

Using Hybrid Model OpenSHMEM + CUDA to Implement the SHOC Benchmark Suite

M. Grodowitz, E. D'Azevedo, S. Powers and N. Imam

Researchers ported several kernels from the Scalable HeterOgeneous Computing (SHOC) benchmark suite from the hybrid MPI + CUDA implementation to an OpenSHMEM + CUDA implementation and demonstrated OpenSHMEM as an alternative communication layer at moderate core counts. This work demonstrates that OpenSHMEM can be effectively used as an alternative to MPI for message passing in inter-device communication in accelerated systems. Specifically, in hybrid models combining OpenSHMEM with CUDA or OpenCL.

This work describes the process of porting the Scalable HeterOgeneous Computing (SHOC) benchmark suite from the hybrid MPI + CUDA implementation to an OpenSHMEM + CUDA implementation. SHOC includes a wide variety of benchmark kernels used to measure accelerator performance in both single node and cluster configurations. The hybrid model implementation attempts to place all major computation on accelerator devices, and uses MPI to synchronize and aggregate results. In some cases, MPI Groups are used to gradually reduce the number of accelerators used for computation as the problem size drops. Porting this behavior to OpenSHMEM required implementing several synchronizing collective operations, and using SHMEM teams to replace MPI Group functionality. Benchmark results on a Cray XK7 system with one GPU per compute node show that SHMEM performance is equal to MPI performance in

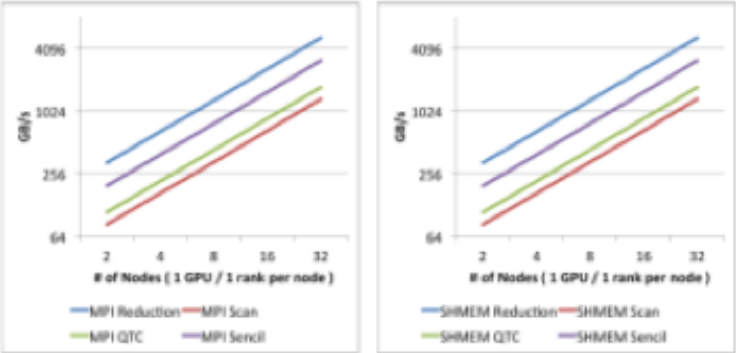

Figure: Truly parallel benchmarks with MPI (left), truly parallel benchmarks with SHMEM (right) indicate similar runtime performance

these hybrid tasks. These results and porting experience show that using OpenSHMEM for accelerator devices benefits from adding functionality for synchronization and teams, and would further benefit from adding support for communication within accelerator kernels.

This work was supported by the DoD, the Extreme Scale Systems Center at ORNL, and the Oak Ridge Leadership Computing Facility

Publication: M. Grodowitz, E. D'Azevedo, S. Powers and N. Imam, "Using Hybrid Model OpenSHMEM + CUDA to Implement the SHOC Benchmark Suite," in Proceedings of Third OpenSHMEM Workshop, Baltimore, MD, USA, August 2016.

How Isotopes Tune Energy Conversion in Polymers

Lei Wang, Jacek Jakowski, Sophya Garashchuk, and Bobby G. Sumpter

This work describes the development and application of a quantum dynamics methodology for multiple quantum nuclei that is suitable for simulations of the isotope effect in conducting polymers. In particular, focus is on an unexpected nuclear quantum effect (NQE) is caused by side-chain H/D substitution in a P3HT/PCBM blend where there is a substantial decrease in the open circuit voltage. The computational approach includes a quadratic basis quantum force correction to account for non-linear contributions of the CH/D anharmonic vibrations. Extensive simulations strongly suggest that isotope-dependent fluctuations may account for the observed NQE by promoting charge transfer via tunneling in P3HT compared to side chain deutrated-P3HT (SD-P3HT)and increasing carrier loss. With this information and the ability to accurately model the isotope effect, specific design of optimally deutrated polymers for OPVs can in principle be achieved.

Work was performed as User Project at the Center for Nanophase Materials Sciences and at the University of South Carolina. This material is based upon work partially supported by the National Science Foundation under Grants No. CHE-1056188 (S.G.). The work regarding P3HT

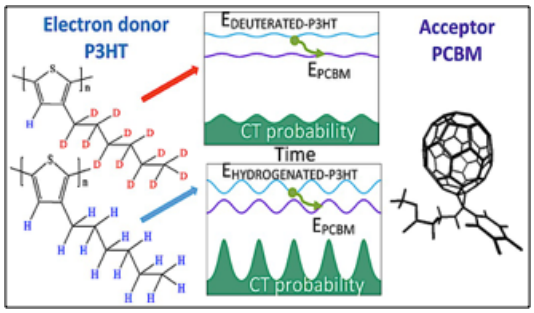

Schematic illustrating the basic mechanism for the isotope effect in side chain deuterated (SD) poly(3-hexylthiophene) (P3HT). The middle panel show the probability of charge transfer to the acceptor molecule (PCBM) and corresponding energy levels; SD-P3HT show shallower fluctuations in probability for charge transfer.

was conducted at the Center for Nanophase Materials Sciences, a U.S. Department of Energy Office of Science User Facility.

Publication: Lei Wang, Jacek Jakowski, Sophya Garashchuk, Bobby G. Sumpter, "Understanding how Isotopes Affect Charge Transfer in P3HT/PCBM: A Quantum Trajectory-Electronic Structure Study with Nonlinear Quantum Corrections", JCTC (2016) DOI: 10.1021/acs.jctc.6b00126

Science Highlights (Continued)

Phase Field Simulations of Autocatalytic Formation of Alpha Lamellar Colonies in Ti-6Al-4V

Bala Radhakrishnan, Gorti Sarma and Suresh Sudharsanam Babu

Researchers demonstrated a possible mechanism for the formation of a lamellar microstructure within layer bands found in additively manufactured titanium alloys using large-scale phase field simulations. Understanding and eliminating the formation of layer bands in structural alloys is essential for reducing the anisotropy in the mechanical properties of additively manufactured components.

We present phase field simulations incorporating contributions due to chemical free energy, anisotropic interfacial energy, and elastic energy due to transformation strain, to demonstrate the nucleation and growth of multiple variants of alpha from undercooled beta in Ti-6Al-4V under isothermal conditions. A new composite nucleation seeding approach is used within the phase field simulations to demonstrate that the presence of a pre-existing strain field can cause the nucleation of specific crystallographic variants of alpha based on minimization of local elastic strain energy. Under conditions where specific combinations of elastic strains exist, for example in the vicinity of one or more pre-existing alpha variants, the nucleation of a new alpha variant is followed by the successive nucleation of the same variant in the form of a lamellar colony by an autocatalytic mechanism. At a given thermodynamic undercooling, the colony structure was favored at a nucleation rate that was low enough to allow sufficient growth of previously nucleated variants before another nucleus formed in their vicinity. Basket weave morphology was formed at higher nucleation rates where multiple nuclei variants grew almost simultaneously under evolving strain fields of several adjacent nuclei.

Research sponsored by the Laboratory Directed Research and Development program at Oak Ridge National Laboratory, managed by UT-Battelle, LLC, under contract DE-AC05- 00OR22725 for the U.S. Department of Energy. This research used resources of the Center for Computational Sciences at Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under contract DE-AC05-00OR22725. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (http://energy.gov/ downloads/ doe-public-access-plan).

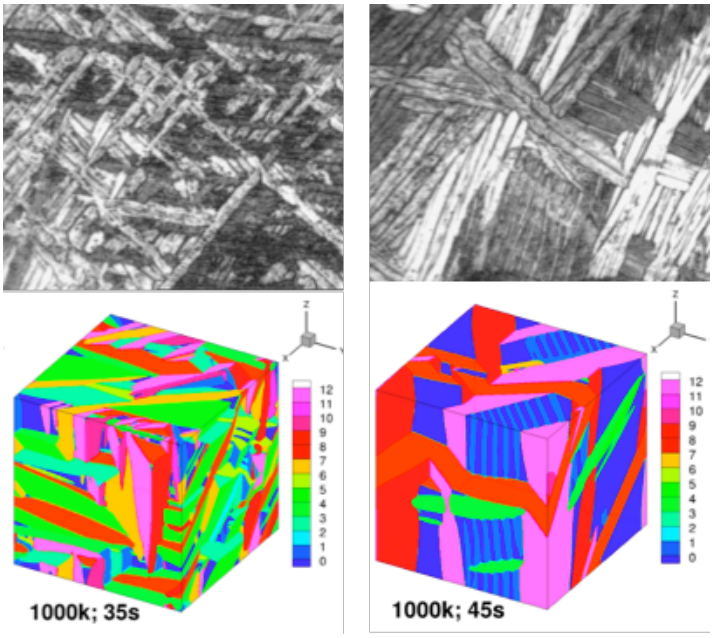

Basket-weave structure (top left) and lamellar structure (top right) found outside and inside the layer bands, respectively, in additively manufactured Ti-6Al-4V alloy components (S.M. Kelly and S.L. Kampe, 2004). Phase field simulations showing the formation of the basket-weave structure (bottom left) and lamellar structure (bottom right) from the present work. The lamellar structure forms by an autocatalytic nucleation and growth mechanism triggered by the presence of well-grown, pre-existing strain fields of nearby crystallographic variants present and at low nucleation rates.

Publication: "Phase Field Simulations of Autocatalytic Formation of Alpha Lamellar Colonies in Ti-6Al-4V"; Bala Radhakrishnan, Sarma Gorti and Suresh Sudharsanam Babu; Metall. Trans. A (2016). DOI:10.1007/s11661-016-3746-6

Science Highlights (Continued)

Designing Functional Molecules to Remove Interactions that Cause Aggregation

Hong-Hai Zhang, Chuanxu Ma, Peter V. Bonnesen, Jiahua Zhu, Bobby G. Sumpter, Jan-Michael Y. Carrillo, Panchao Yin, Yangyang Wang, An-Ping Li, and Kunlun Hong

For the first time, researchers were able to synthesize Poly(5-alkyl-2,3-thiophene)s (P5AT), an ortho-connected isomer of poly(3-alkyl-2,5-thiophene)s (P3AT). P5ATs show distinct helical structures with different optoelectronic properties compared to their P3AT isomeric counterparts and are representative of a new class of potential optoelectronic materials.

Conjugated polyarylenes have found a wide range of applications including optical sensors, electrochromic devices, field effect transistors, and organic solar cells. However often the interactions driven by the conjugated pi-backbone causes difficulties in solution processing due to strong aggregation of the polymer chains. One approach to uncouple these backbone interactions between polymer chains to mediate aggregation is to introduce a helix geometry for the backbone. This could be done by using ortho-connected isomers, poly(5-alkyl-2,3-thiophene)s (P5AT) but such polyarylenes have never been reported because of the difficulty in their syntheses. This work marks the first synthesis of regioregular P5AT; achieved via controlled Suzuki cross-coupling polymerization with PEPPSI-IPr as catalyst, affording the polymers with tunable molecular weight, narrow polydispersity (PDI), and well-defined functional end groups at the gram scale. The helical geometry of the P5ATs predicted by computational modeling was confirmed by a combination of NMR, small-angle X-ray scattering (SAXS), and scanning tunneling microscopy (STM). Since the Suzuki cross-coupling polymerization tolerates a wide range of functionalized

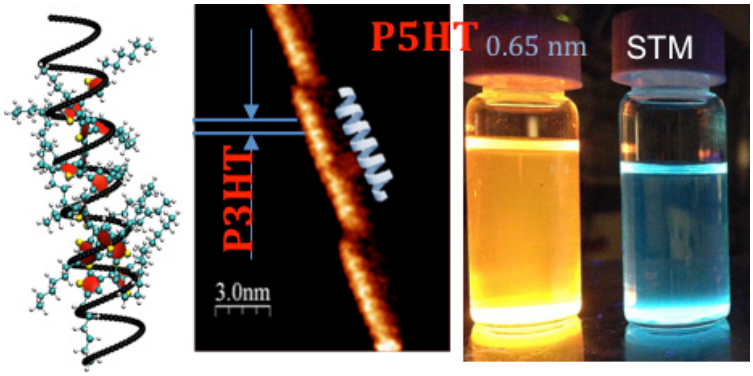

Left: Computational modeling predicted structure for P5BT. Middle: STM image of P5BT on HOPG showing a 0.65 nanometer helical pitch. Right: Photographs of the THF solutions of P3HT and P5HT under 360 nm UV light showing that P5HT has green photoluminescence under irradiation.

aromatic monomers, this method can be applicable for the synthesis of poly(o-arylene)s with aromatic rings other than thiophene and thus opens access to a new class of poyarylenes where the combination of helical structure, conjugated backbone, and photoluminescence properties make P5ATs good candidates for a range of applications in optoelectronics and related field applications.

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

Publication: "Helical Poly(5-alkyl-2,3-thiophene)s: Controlled Synthesis and Structure Characterization"; Hong-Hai Zhang, Chuanxu Ma, Peter V. Bonnesen, Jiahua Zhu, Bobby G. Sumpter, Jan-Michael Y. Carrillo, Panchao Yin, Yangyang Wang, An-Ping Li, and Kunlun Hong

News

OLCF Researchers Scale R to Tackle Big Science Data Sets

Sometimes lost in the discussion around big data is the fact that big science has long generated huge data sets. "In fact, large-scale simulations that run on leadership-class supercomputers work at such high speeds and resolution that they generate unprecedented amounts of data. The size of these datasets—ranging from a few gigabytes to hundreds of terabytes—makes managing and analyzing the resulting information a challenge in its own right," notes an article posted on Oak Ridge National Laboratory site yesterday.

Now, a group at ORNL's Oak Ridge Leadership Computing Facility has developed an updated version of R – programming with big data in R (pbdR) – that is helping OLCF users cut large data sets down to size.

Read more here - https://www.hpcwire.com/2016/07/06/olcf-researchers-scale-r-tackle-big-science-data-sets/

Emilio Ramirez Sees Better Outcomes

Just a few years ago, Emilio Ramirez spent his days operating and adjusting settings to optimize thermal performance at a Central California bioenergy power plant.

Emilio, a California native who is now a University of Tennessee doctoral candidate working with Oak Ridge National Laboratory, acquired his first experience with bioenergy as the lead operator of a biomass circulating fluidized-bed combustion facility. Although the work was steady and intriguing, Emilio knew he wanted to contribute to bioenergy on a national scale.

Read the full story here - https://www.ornl.gov/news/emilio-ramirez-path-biomass-research-full-twists-turns-and-personal-discovery.

New Genomics Pipeline Combines AWS, Local HPC, and Supercomputing

Declining DNA sequencing costs and the rush to do whole genome sequencing (WGS) of large cohort populations – think 5000 subjects now, but many more thousands soon – presents a formidable computational challenge to researchers attempting to make sense of large cohort datasets. No single architecture is best. This month researchers report developing a hybrid approach that combines cloud (AWS), local high performance compute (LHPC) clusters, and supercomputers.

Their fascinating paper, A hybrid computational strategy to address WGS variant analysis in >5000 samples, spells out in some detail the obstacles associated with using each resource and how to divide the work to maximize throughput and minimize cost. Computational resources used included: Amazon AWS; a 4000-core in-house cluster at Baylor College of Medicine; IBM power PC Blue BioU at Rice University and Rhea at Oak Ridge National Laboratory (ORNL). DNAnexus was also a collaborator.

Read the full story here - https://www.hpcwire.com/2016/09/22/new-genomics-pipeline-combines-aws-local-hpc-supercomputing/.

Community Service

- Rizwan A. Ashraf, program committee, 2nd International Workshop on Fault Tolerant Systems
- David E. Bernholdt, program committee, Hawaii International Conference on System Science 2017 (HICCS-50)
- David E. Bernholdt, peer reviewer, Journal of Computational Science (JOCS)
- Michael Brim, peer reviewer, Journal Parallel Computing
- Christian Engelmann, peer reviewer, Journal of Computational Science (JOCS)
- Saurabh Hukerikar, peer reviewer, Journal of Supercomputing
- Saurabh Hukerikar, peer reviewer, Journal of Supercomputing [Special Issue on Emerging Platform Technologies]
- Jungwon Kim, External Reviewer, ASPLOS 2017: The 22nd ACM International Conference on Architectural Support for Programming Languages and Operating Systems.
- Seyong Lee, Program Committee, ICPADS 2016: The 22nd IEEE International Conference on Parallel and Distributed Systems
- Seyong Lee, Program Committee, WACCPD 2016: Third Workshop on Accelerator Programming Using Directives
- Seyong Lee, Program Committee, LHAM 2016: International Workshop on Legacy HPC Application Migration
- Tonglin Li, peer reviewer, IEEE Access
- Tonglin Li, peer reviewer, IEEE Transactions on

Community Service (Continued)

Industrial Informatics

- Tonglin Li, program committee, 3rd IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (IEEE/ACM BDCAT 2016)
- Sarat Sreepathi, Reviewer for Journal of Applied Soft Computing.
- Sarat Sreepathi, Reviewer for Journal of Water and Climate Change.
- Jeffrey Vetter, Program Committee, PPOPP, 2017
- Jeffrey Vetter, Program Committee, ISC, 2017
- Jeffrey Vetter, Program Committee, Exascale Requirements Review for Advanced Scientific Computing Research (ASCR), 2016
- Jeffrey Vetter, Technical Papers Program Committee, SC16

Events

OpenSHMEM 2016: Third workshop on OpenSHMEM and Related Technologies

OpenSHMEM 2016 is the third event in the OpenSHMEM and Related Technologies workshop series. The workshop was organized by Oak Ridge National Laboratory and held in Baltimore, Maryland, USA, and it was and sponsored by ORNL, DoD, Cray, Mellanox, Intel, and SGI.

The workshop had several HPC pioneers speaking at the workshop, which included Steve Oberlin (NVIDIA CTO), James Sexton (IBM Fellow), and DK Panda (Professor at Ohio State U.). In addition to keynote talks, the first two days of the workshop had over 20 technical talks, which included paper talks, and vendor talks.

The third day of the workshop was focused on developing the OpenSHMEM specification. This was another turning point for the OpenSHMEM community, since we have established the procedures for OpenSHMEM committee operations. This should streamline the process of bringing in new changes to the specification. Steve Poole (DoD) was elected as OpenSHMEM Committee chair and Manjunath Gorentla Venkata (ORNL) was elected as OpenSHMEM Committee secretary.

9th MPM Workshop

John Turner and Stuart Slattery attended the 9th MPM Workshop on September 8th and 9th in Portland, Oregon. They presented ongoing work at ORNL in developing scalable material point method (MPM) algorithms for next-generation supercomputers and discussed ORNL work in the modeling of electric vehicle batteries and additive manufacturing. More information on the meeting can be found at (http://www.mesomechanics.com/home.html) with meeting presentation slides available upon request.

CAM Seminar Series

- Dr. Jeremy Trageser, George Washington University - A Fourth-order Nonlocal Operator and Its Connection with Its Local Analogue
- Dr. Erik von Schwerin, University of Delaware - Optimal Mesh Hierarchies in Multilevel Monte Carlo Methods
- Dr. Kody Law, Oak Ridge National Laboratory - Multilevel Monte Carlo for Inference
- Dr. Miroslav Stoyanov, Oak Ridge National Laboratory - A Reduced Basis Multi-Level Method for High Dimensional Stochastic Systems
- Ms. Liu Liu, University of Wisconsin - Madison - An Asymptotic-preserving Stochastic Galerkin Method for the Semiconductor Boltzmann Equation with Random Inputs and Diffusive Scalings
- Dr. Shi Jin, University of Wisconsin - Madison - Uncertainty Quantification for Multiscale Kinetic Equations with Uncertain Coefficients
- 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

CSMD Seminars

- CJ Newburn: A Path Forward to Scaled Execution
- Dr. Mehmet Belviranli: Increasing Resource Utilization in Heterogeneous Architectures
- Dr. Jeremiah Willcock: Abstractions for Parallel Graph Algorithms
- Dr. Joel E. Denny: NVL-C: Static Analysis Techniques for Efficient, Correct Programming of Non-Volatile Main Memory Systems
- Professor Hank Childs: In Situ at the Exascale: Motivations, Instantiations, and Opportunities
- Dr.-Ing. Bernd Mohr: Multicore Performance Analysis at Scale: From Single Node to one Million HPC Cores

Events (Continued)

- Dr. Fabien Delalondre: Challenges and possible workflow, infrastructure and computing solutions to support multilevel brain modeling within the next decade
- Dr. Jerry Chow and Dr. Jay Gambetta: Quantum error detection, high-fidelity control, and experiencing other things quantum
- Camila A. Ramirez: Implementation of SIR particle filter for radiation source detection
- Dr. Andrey Prokopenko: Developing an algebraic multigrid software package to help applications
- Dr. Peter M. Kogge: Introducing Emu Solutions' Migratory Thread and Memory-Side Processing Technology
- Dr. Faisal Shah Khan: Non-cooperative Games and Quantum Computing
- Dr. Ryan Bond: Extension of Kestrel to General Thermochemical Models

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.

Travis Humble and Raphael C. Pooser's received a patent for "Self-Correcting Random Number Generator"

The patent's abstract describes a system and method for generating random numbers. The system may include a random number generator (RNG), such as a quantum random number generator (QRNG) configured to self-correct or adapt in order to substantially achieve randomness from the output of the RNG. By adapting, the RNG may generate a random number that may be considered random regardless of whether the random number itself is tested as such. As an example, the RNG may include components to monitor one or more characteristics of the RNG during operation, and may use the monitored characteristics as a basis for adapting, or self-correcting, to provide a random number according to one or more performance criteria.

Publications of Note

A three-dimensional meso-macroscopic model for Li-Ion intercalation batteries

S. Allu, S. Kalnaus, S. Simunovic, J. Nanda, J.A. Turner, S. Pannala; Journal of Power Sources, 325 (2016): 42-50.

In this paper we present a three-dimensional computational formulation for electrode-electrolyte-electrode system of Li-Ion batteries. The physical consistency between electrical, thermal and chemical equations is enforced at each time increment by driving the residual of the resulting coupled system of nonlinear equations to zero. The formulation utilizes a rigorous volume averaging approach typical of multiphase formulations used in other fields and recently extended to modeling of supercapacitors. Unlike existing battery modeling methods which use segregated solution of conservation equations and idealized geometries, our unified approach can model arbitrary battery and electrode configurations. The consistency of multi-physics solution also allows for consideration of a wide array of initial conditions and load cases. The formulation accounts for spatio-temporal variations of material and state properties such as electrode/void volume fractions and anisotropic conductivities. The governing differential equations are discretized using the finite element method and solved using a nonlinearly consistent approach that provides robust stability and convergence. The new formulation was validated for standard Li-ion cells and compared against experiments. Its scope and ability to capture spatio-temporal variations of potential and lithium distribution is demonstrated on a prototypical three-dimensional electrode problem.

Publications of Note (Continued)

Phase Field Simulations of Autocatalytic Formation of Alpha Lamellar Colonies in Ti-6Al-4V

B. Radhakrishnan, S.B. Gorti and S.S. Babu, Metall. Trans. A, 10.1007/s11661-016-3746-6, Sep. 2016.

We present phase field simulations incorporating energy contributions due to thermodynamics, and anisotropic interfacial and strain energies, to demonstrate the nucleation and growth of multiple variants of alpha from beta in Ti-6Al-4V under isothermal conditions. The simulations focused on the effect of thermodynamic driving force and nucleation rate on the morphology of the transformed alpha assuming that the partitioning of V between beta and alpha is negligible for short isothermal holds. The results indicate that a high nucleation rate favors the formation of the basket-weave structure. However, at a lower nucleation rate the simulations show the intragranular nucleation of a colony structure by an autocatalytic nucleation mechanism adjacent to a pre-existing alpha variant. New side-plates of the same variant appear to nucleate progressively and grow to form the colony. The isothermal simulation results are used to offer a possible explanation for the transition from a largely basket weave structure to a colony structure inside narrow layer bands occurring during continuous heating and cooling conditions encountered during laser additive manufacturing of Ti-6Al-4V.

Understanding How Isotopes Affect Charge Transfer in P3HT/PCBM: A Quantum Trajectory-Electronic Structure Study with Nonlinear Quantum Corrections

Lei Wang, Jacek Jakowski, Sophya Garashchuk, and Bobby G. Sumpter

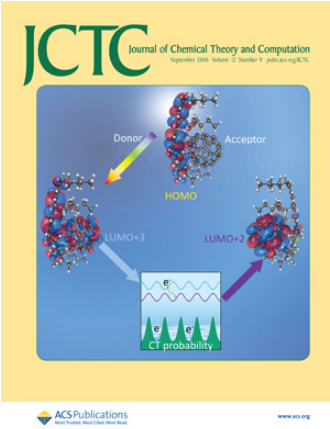

The experimentally observed effect of selective deuterium substitution on the open circuit voltage for a blend of poly(3-hexylthiophene) (P3HT) and [6,6]-phenyl-C61-butyric acid methyl ester (PCBM; Nat. Commun. 2014, 5, 3180) is explored using a 221-atom model of a polymer-wrapped PCBM molecule. The protonic and deuteronic wave functions for the H/D isotopologues of the hexyl side chains are described within a quantum trajectory/electronic structure approach where the dynamics is performed with newly developed nonlinear corrections to the quantum forces, necessary to describe the nuclear wave functions; the classical forces are generated with a density functional tight binding method. The resulting protonic and deuteronic time-dependent wave functions are used to assess the effects of isotopic substitution (deuteration) on the energy gaps relevant to the charge transfer for the donor and acceptor electronic states. While the isotope effect on the electronic energy levels is found negligible, the quantum-induced fluctuations of the energy gap between the charge transfer and charge separated states due to nuclear wave functions may account for experimental trends by promoting charge transfer in P3HT:PCBM and increasing charge recombination on the donor in the deuterium substituted P3HT:PCBM.

Self-Assembly: Thermoreversible Morphology and Conductivity of a Conjugated Polymer Network Embedded in Block Copolymer Self-Assemblies (Small 35/2016)

Youngkyu Han, Jan-Michael Y. Carrillo, Zhe Zhang, Yunchao Li, Kunlun Hong, Bobby G. Sumpter, Michael Ohl, Mariappan Parans Paranthaman, Gregory S. Smith, Changwoo Do

The thermo-responsive structural and electrical behaviors of polymeric self-assemblies comprising of PEO-PPO-PEO block copolymers and poly(thiophene) derivatives are investigated by C. Do and co-workers on page 4857. A synergetic usage of small-angle neutron and x-ray scattering with computer simulations reveals that a tunable morphology of triblock copolymer templates enables to thermo-reversibly control distribution and structure of the embedded conjugated polymers, strongly coupled to unique conduction property.

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