ISSUE
21
SUMMER
2016

QUARTERLY NEWSLETTER FOR THE COMPUTER SCIENCE AND MATHEMATICS DIVISION

# CSMD

this issue

Science Highlights P.01
News P.18
Community Service P.19
Events P.20
Software P.22
Awards and Recognition P.23
Publications of Note P.23
About CSMD P.26

## Science Highlights

### Improving Predictions of Heterotrophic Respiration

*B. Bond-Lamberty, D. Epron, J. Harden, M. E. Harmon, F. M. Hoffman, J. Kumar, A. D. McGuire, and R. Vargas*

Researchers proposed improved representation of heterotrophic respiration (HR) in Earth system models by grouping metabolism and flux characteristics across space and time.

Heterotrophic respiration (HR), the aerobic and anaerobic processes mineralizing organic matter, is a key carbon flux but one impossible to measure at scales significantly larger than small experimental plots. This impedes our ability to understand carbon and nutrient cycles, benchmark models, or reliably upscale point measurements. Given that a new generation of highly mechanistic, genomic-specific global models is not imminent, we suggest that a useful step to improve this situation is the development of Decomposition Functional Types (DFTs). Analogous to plant functional types (PFTs), DFTs would abstract and capture important differences in HR metabolism and flux dynamics, allowing modelers and experimentalists to efficiently group and vary these characteristics across space and time. We applied cluster analysis to show how annual HR can be broken into distinct groups associated with global variability in biotic and abiotic factors, and we demonstrated that these groups are distinct from, but complementary to, PFTs. In this position paper, we suggested priorities for next steps to build a foundation for DFTs in global models to provide the ecological and climate change communities with robust, scalable estimates of HR.

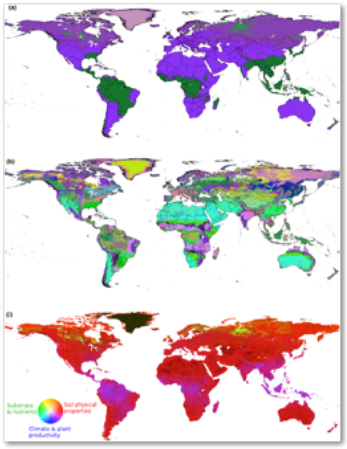

***Example cluster analyses delineating DFTs from 11 global climatic, edaphic, carbon flux, and topographic characteristics.***

This research is the product of a working group on heterotrophic respiration led by Mark Harmon and sponsored by the National Science Foundation, which funded meeting and travel expenses. Contributions of Bond-Lamberty, Hoffman, and Kumar were supported by the DOE Office of Science, BER, Regional & Global Climate Modeling and Terrestrial Ecosystem Science Programs in the Climate & Environmental Sciences Division.

Publication: B. Bond-Lamberty, D. Epron, J. Harden, M. E. Harmon, F. M. Hoffman, J. Kumar, A. D. McGuire, and R. Vargas, "Estimating heterotrophic respiration at large scales: Challenges, approaches, and next steps." Ecosphere 7, (2016). doi:10.1002/ecs2.1380.

OAK RIDGE National Laboratory

# Science Highlights (Continued)

## NVL-C: Static Analysis Techniques for Efficient, Correct Programming of Non-Volatile Main Memory Systems

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

A novel programming system that extends C with intuitive, language-level support for programming Non-Volatile Memory (NVM) as persistent, high-performance main memory.

Computer architecture experts expect that NVM hierarchies will play a more significant role in future systems including mobile, enterprise, and HPC architectures. With this expectation in mind, we present NVL-C: a novel programming system that facilitates the efficient and correct programming of NVM main memory systems. The NVL-C programming abstraction extends C with a small set of intuitive language features that target NVM main memory, and can be combined directly with traditional C memory model features for DRAM. We have designed these new features to enable compiler analyses and run-time checks that can improve performance and guard against a number of subtle programming errors, which, when left uncorrected, can corrupt NVM-stored data. Moreover, to enable recovery of data across application or system failures, these NVL-C features include a flexible directive for specifying NVM transactions. So that our implementation might be extended to other compiler front ends and languages, the majority of our compiler analyses are implemented in an extended version of LLVM's intermediate representation (LLVM IR). We evaluate NVL-C on a number of applications to show its flexibility, performance, and correctness.

This material is based upon work supported by the U.S. Department of Energy (DOE), Office of Science, Office of Advanced Scientific Computing Research. This manuscript has been authored by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725 with the DOE. The United States Government (USG) retains and the publisher, by accepting the article for publication, acknowledges that the USG retains a non-exclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for USG purposes. The DOE 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).

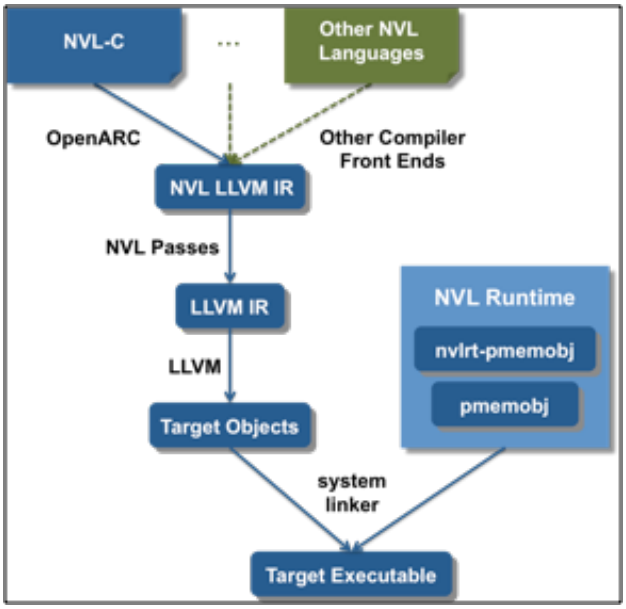


***The NVL-C compiler front end is an extension of the OpenARC compiler framework. So that the NVL-C compiler implementation might be extended to other compiler front ends and languages, the majority of the compiler analyses are implemented as an extension of the popular LLVM compiler infrastructure. NVL-C's runtime currently uses Intel's pmem library for low-level NVM memory management and transactions. However, the interface of the runtime is encapsulated so that pmem can easily be replaced by alternate runtime implementations in the future without adjusting the compiler.***

Publication: "*NVL-C: Static Analysis Techniques for Efficient, Correct Programming of Non-Volatile Main Memory Systems*"; Joel E. Denny, Seyong Lee, and Jeffrey S. Vetter; HPDC (2016). DOI:10.1145/2907294.2907303

# Science Highlights (Continued)

## Human-induced Greening of the Northern Extratropical Land Surface

*Jiafu Mao, Aurélien Ribes, Binyan Yan, Xiaoying Shi, Peter E. Thornton, Roland Séférian, Philippe Ciais, Ranga B. Myneni, Hervé Douville, Shilong Piao, Zaichun Zhu, Robert E. Dickinson, Yongjiu Dai, Daniel M. Ricciuto, Mingzhou Jin, Forrest M. Hoffman, Bin Wang, Mengtian Huang, and Xu Lian*

Two long-term satellite leaf area index (LAI) datasets, 19 coupled earth system models (ESMs) and a formal detection and attribution algorithm were used to attribute changes of vegetation activity in the northern-extratropical latitudes (NEL) for the period 1982-2011.

Significant land greening in the NEL has been documented through satellite observations during the past three decades. This enhanced vegetation growth has broad implications for surface energy, water and carbon budgets, and ecosystem services across multiple scales. Discernable human impacts on the Earth's climate system have been revealed by using statistical frameworks of detection and attribution. These impacts, however, were not previously identified on the NEL greening signal, due to the lack of long-term observational records, possible bias of satellite data, different algorithms used to calculate vegetation greenness, and the lack of suitable simulations from coupled Earth system models (ESMs). Here we have overcome these challenges in order to attribute recent changes in NEL vegetation activity. We used two 30-year-long remote-sensing-based LAI datasets, simulations from 19 coupled ESMs with interactive vegetation, and a formal detection and attribution algorithm. Our findings reveal that the observed greening record is consistent with an assumption of anthropogenic forcings, where greenhouse gases play a dominant role, but is not consistent with simulations that include only natural forcings and internal climate variability. Given the strong evidence provided here of historical human-induced greening in the northern extratropics, society should consider both intended and unintended consequences of its interactions with terrestrial ecosystems and the climate system.

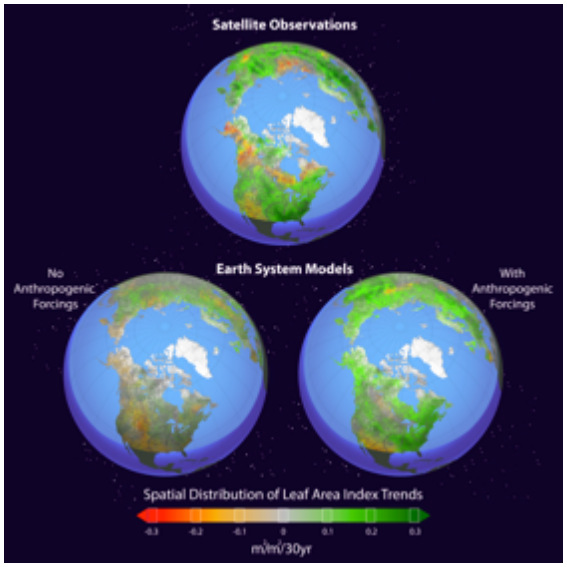


***Spatial distribution of LAI trends observed by satellite and simulated by CMIP5 models over the period 1982–2011.***

J. Mao, X. Shi, P.E. Thornton, D.M. Ricciuto, and F.M. Hoffman are supported by DOE Office of Science, Biological and Environmental Research, including support from the following programs: Regional and Global Climate Modeling Program (ORNL BGC-Feedbacks SFA), Terrestrial Ecosystem Science Program (ORNL TES SFA), and Earth System Modeling (ACME project)

Publication: Jiafu Mao*, Aurélien Ribes, Binyan Yan, Xiaoying Shi, Peter E. Thornton, Roland Séférian, Philippe Ciais, Ranga B. Myneni, Hervé Douville, Shilong Piao, Zaichun Zhu, Robert E. Dickinson, Yongjiu Dai, Daniel M. Ricciuto, Mingzhou Jin, Forrest M. Hoffman, Bin Wang, Mengtian Huang, and Xu Lian, 2016. Human-induced greening of the northern extratropical land surface. Nature Climate Change, 10.1038/nclimate3056.

# Science Highlights (Continued)

## Modeling the Carbon Cycle as a Nonautonomous System

*NIMBioS Working Group on Nonautonomous Systems and the Terrestrial Carbon Cycle (see http://www.nimbios.org/workinggroups/WG_ccycle)*

Researchers developed an understanding of transit time and mean age dynamics of terrestrial carbon storage with time-dependent parameters and inputs.

We develop a theory for transit times and mean ages for nonautonomous compartmental systems. Using the McKendrick–von Förster equation, we show that the mean ages of mass in a compartmental system satisfy a linear nonautonomous ordinary differential equation that is exponentially stable. We then define a nonautonomous version of transit time as the mean age of mass leaving the compartmental system at a particular time and show that our nonautonomous theory generalises the autonomous case. We apply these results to study a nine-dimensional nonautonomous compartmental system modeling the terrestrial carbon cycle, which is a modification of the Carnegie–Ames–Stanford approach model, and we demonstrate that the nonautonomous versions of transit time and mean age differ significantly from the autonomous quantities when calculated for that model.

This research is the result of a National Institute for Mathematical and Biological Synthesis (NIMBioS) working group on “Nonautonomous linear system of the terrestrial carbon cycle: Mathematical and ecological properties and their uses in guiding carbon research” led by Yiqi Luo, Maria Leite, and Ying Wang (see http://www.nimbios.org/workinggroups/WG_ccycle) that was funded by

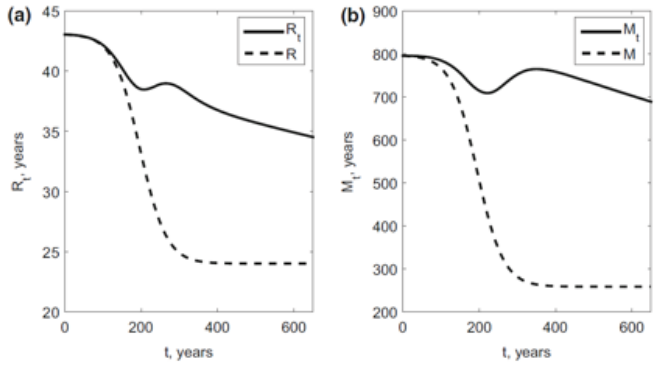


***Using our nonautonomous theory for representing a nine-pool terrestrial carbon cycle model, we showed an order of magnitude difference in the absolute values of mean transit time, Rt, and mean age, Mt. Moreover, significant differences were shown between these nonautonomous properties and the instantaneous quantities, R and M, which represent the autonomous model.***

the National Science Foundation (NSF), which paid for travel and meeting expenses. Hoffman's contribution was supported by the DOE Office of Science, Biological and Environmental Research (BER), Regional and Global Climate Modeling (RGCM) Program in the Climate and Environmental Sciences Division (CESD).

Publication: Rasmussen, M., A. Hastings, M. J. Smith, F. B. Agusto, B. M. Chen-Charpentier, F. M. Hoffman, J. Jiang, K. E. O. Todd-Brown, Y. Wang, Y.-P. Wang, and Y. Luo (2016), Transit times and mean ages for nonautonomous and autonomous compartmental systems, J. Math. Biol., doi:10.1007/s00285-016-0990-8.

# Science Highlights (Continued)

### Phosphorus feedbacks may constrain tropical ecosystem responses to changes in atmospheric $CO_2$ and climate

*X. Yang, P.E. Thornton, D.M. Ricciuto, and F.M. Hoffman*

This study shows that the coupling of P cycle in land surface model results in more realistic spatial pattern of simulated ecosystem productivity in the Amazon region. Through exploratory simulations this study points to the need for more tropical field measurements under different temperature/humidity conditions with different soil P availability.

It is being increasingly recognized that carbon-nutrient interactions play important roles in regulating terrestrial carbon cycle responses to increasing $CO_2$ in the atmosphere and climate change. Nitrogen-enabled models in CMIP5 showed that accounting for nitrogen greatly reduces the negative feedback between land ecosystems and atmospheric $CO_2$. None of the CMIP5 models has considered P as a limiting nutrient, although P has been considered the most limiting nutrient in lowland tropical forests. In this study, scientists from Oak Ridge National Lab investigated the effects of P availability on carbon cycling in the Amazon region using a P-enabled land surface model. Model simulations demonstrate that $CO_2$ fertilization effect in the Amazon region may be greatly overestimated if P cycling were not considered. Exploratory simulations highlighted the importance of considering the interactions between carbon, water, and nutrient cycling (both nitrogen and phosphorus) for the prediction of future carbon uptake in tropical ecosystems.

X. Yang, P.E. Thornton, D.M. Ricciuto, and F.M. Hoffman are supported by DOE Office of Science, Biological and Environmental Research, including support from the following programs: Regional and Global Climate Modeling Program (ORNL BGC-Feedbacks SFA), Terrestrial Ecosystem Science Program (ORNL TES SFA and NGEE Tropics), Earth System Modeling (ACME project).

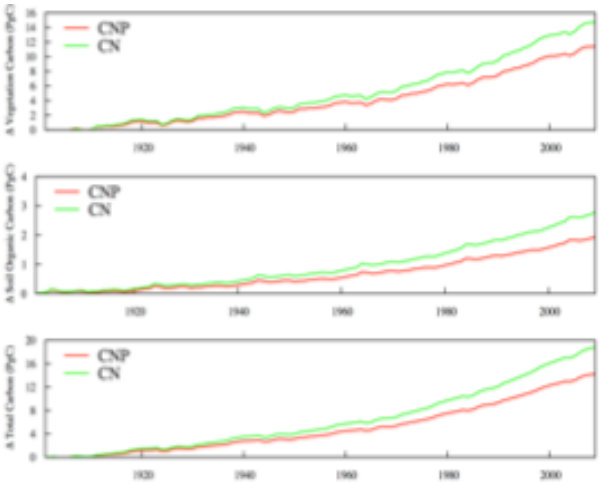


***This graphic shows the changes in carbon stocks in vegetation, soil, total(vegetation plus soil) for the period of 1900-2009 associated with historical changes in atmospheric $CO_2$. The CNP simulated carbon accumulation is about 26% lower than that simulated by the CN model, which is mainly due to low soil P availability in most of the Amazon region that causes a much weaker response of plant growth to increasing atmospheric $CO_2$.***

Publication: X.Yang, P. E. Thornton, D. M. Ricciuto, and F.M. Hoffman, “Phosphorus feedbacks may constrain tropical ecosystem responses to changes in atmospheric CO2 and climate”, Geophysical Research Letters (accepted, 2016)

# Science Highlights (Continued)

## Responses of Two Nonlinear Microbial Models to Warming and Increased Carbon Input

*NIMBioS Working Group on Nonautonomous Systems and the Terrestrial Carbon Cycle (see http://www.nimbios.org/workinggroups/WG_ccycle)*

Researchers improved understanding of significance of soil microbial processes in soil carbon responses to climate change.

A number of nonlinear microbial models of soil carbon decomposition have been developed. Some of them have been applied globally but have yet to be shown to realistically represent soil carbon dynamics in the field. A thorough analysis of their key differences is needed to inform future model developments. Here we compare two nonlinear microbial models of soil carbon decomposition: one based on reverse Michaelis–Menten kinetics (model A) and the other on regular Michaelis–Menten kinetics (model B). Using analytic approximations and numerical solutions, we find that the oscillatory responses of carbon pools to a small perturbation in their initial pool sizes dampen faster in model A than in model B. Soil warming always decreases carbon storage in model A, but in model B it predominantly decreases carbon storage in cool regions and increases carbon storage in warm regions. For both models, the $CO_2$ efflux from soil carbon decomposition reaches a maximum value some time after increased carbon input (as in priming experiments). This maximum $CO_2$ efflux (Fmax) decreases with an increase in soil temperature in both models. However, the sensitivity of Fmax to the increased amount of carbon input increases with soil temperature in model A but decreases monotonically with an increase in soil temperature in model B. These differences in the responses to soil warming and carbon input between the two nonlinear models can be used to discern which model is more realistic when compared to results from field or laboratory experiments. These insights will contribute to an improved understanding of the significance of soil microbial processes in soil carbon responses to future climate change.

This research is the result of a National Institute for Mathematical and Biological Synthesis (NIMBioS) working group on “Nonautonomous linear system of the terrestrial carbon cycle: Mathematical and ecological properties and their uses in guiding carbon research” led by Yiqi Luo, Maria Leite, and Ying Wang (see http://www.nimbios.org/workinggroups/WG_ccycle) that was funded by the National Science Foundation (NSF), which paid for travel and meeting expenses. Hoffman's contribution was supported by the DOE Office of Science, Biological and Environmental Research (BER), Regional and Global Climate Modeling (RGCM) Program in the Climate and Environmental Sciences Division (CESD).

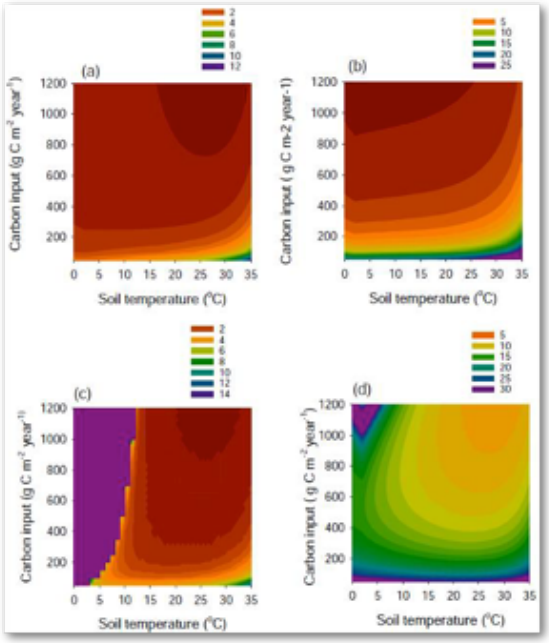


***Half-life (a, b) and period (c, d) in units of year for model A (a, c) and B (b, d). The purple region represents non-oscillatory region for model A in (c), and a period greater than 30 years for model B in (d).***

Publication: Wang, Y. P., J. Jiang, B. Chen-Charpentier, F. B. Agusto, A. Hastings, Forrest M. Hoffman, M. Rasmussen, M. J. Smith, K. Todd-Brown, Y. Wang, X. Xu, and Y. Q. Luo (2016), Responses of two nonlinear microbial models to warming and increased carbon input, Biogeosci., 13(4):887–902. doi:10.5194/bg-13-887-2016.

# Science Highlights (Continued)

## IMPACC: A Tightly Integrated MPI+OpenACC Framework Exploiting Shared Memory Parallelism

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

IMPACC is a MPI+OpenACC framework for heterogeneous clusters. IMPACC starts with the observation that current MPI+OpenACC model has some inefficiencies and complexities to orchestrate two orthogonal programming models. IMPACC solves these problems by exploiting shared memory parallelism and the tight integration of MPI and OpenACC.

In this work, we propose IMPACC, an MPI+OpenACC framework for heterogeneous accelerator clusters. IMPACC tightly integrates MPI and OpenACC, while exploiting the shared memory parallelism in the target system. IMPACC dynamically adapts the input MPI+OpenACC applications on the target heterogeneous accelerator clusters to fully exploit target system-specific features. IMPACC provides the programmers with the unified virtual address space, automatic NUMA-friendly task-device mapping, efficient integrated communication routines, seamless streamlining of asynchronous executions, and transparent memory sharing. We have implemented IMPACC and evaluated its performance using three heterogeneous accelerator systems, including Titan supercomputer. Results show that IMPACC can achieve easier programming, higher performance, and better scalability than the current MPI+OpenACC model.

This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

This material is based upon work supported by the National Science Foundation under Grant Number 1137097 and by the University of Tennessee through the Beacon Project. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or the University of Tennessee. The authors would like to thank NVIDIA for providing access to their PSG Cluster. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research. 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 non-exclusive, paid-up, irrevocable, world-wide 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).

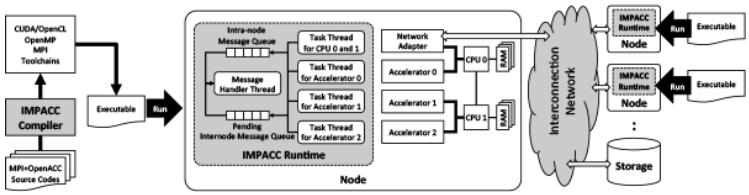


***The IMPACC framework consists of two parts, the IMPACC compiler and the IMPACC runtime. The IMPACC compiler translates the compute-intensive codes augmented by parallel or kernels construct in the input MPI+OpenACC source codes into accelerator kernel codes written in CUDA C and OpenCL C. The IMPACC runtime in every node launches the same number of program instances as the number of available accelerator devices in the node. The runtime implements a task as a lightweight user-level thread and assigns a distinct accelerator to each task.***

Publication: : "IMPACC: A Tightly Integrated MPI+OpenACC Framework Exploiting Shared Memory Parallelism"; Jungwon Kim, Seyong Lee, and Jeffrey S. Vetter; the 25th ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC16). DOI: 10.1145/2907294.2907302

# Science Highlights (Continued)

## SnuCL-D: A Distributed OpenCL Framework using Redundant Computation and Data Replication

*Junghyun Kim, Gangwon Jo, Jaehoon Jung, Jungwon Kim, and Jaejin Lee*

SnuCL-D is the first work to show scalability of OpenCL application written for a single node on large-scale systems. This work was published in PLDI16. PLDI is the premier forum for researchers, developers, practitioners, and students to present research on programming language design and implementation.

In this work, we propose a scalable and distributed OpenCL framework called SnuCL-D for large-scale clusters. SnuCL-D overcomes the performance bottleneck incurred by four different types of overheads in previous centralized approaches, namely command-queueing, command-scheduling, command-delivery, and host-data-transfer overheads. The experimental results show that SnuCL-D scales much better than SnuCL for applications that execute a large number of commands. SnuCL-D is more than 20 times faster than SnuCL on 512 nodes for these applications. They also indicate that SnuCL-D and MPI-Fortran are comparable in performance for most of the applications up to 128 nodes. For 512 nodes, MPI-Fortran outperforms SnuCL-D because the amount of work in an OpenCL kernel is too small to amortize the inherent overhead of the vendor-specific OpenCL runtime. We expect that SnuCL-D and MPI-Fortran are comparable in performance for more than 128 nodes with much bigger input sizes. SnuCL-D efficiently and transparently extends the OpenCL programming model to clusters. Neither modifying the original code nor using any communication library is necessary to use SnuCL-D. An OpenCL application written for multiple devices can be executed efficiently

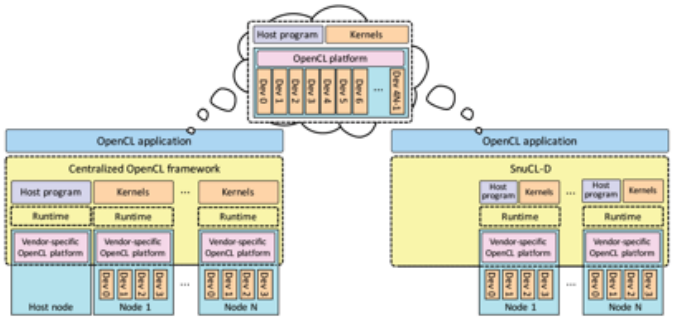


***All previous approaches (left-hand side) have a centralized host node that executes the OpenCL host program. Other nodes are compute nodes and perform kernel computations on their compute devices. These centralized approaches introduce significantly communication overhead on large-scale systems. The totally distributed redundant computation and data replication makes SnuCL-D (right-hand side) outperforms the centralized approach.***

on a large-scale cluster as a whole. The source code of SnuCL-D is publicly available at the URL http://aces.snu.ac.kr.

This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIP) (No. 2013R1A3A2003664). ICT at Seoul National University provided research facilities for this study.

Publication: “A Distributed OpenCL Framework using Redundant Computation and Data Replication”; Junghyun Kim, Gangwon Jo, Jaehoon Jung, Jungwon Kim, and Jaejin Lee; the 37th ACM SIGPLAN conference on Programming Language Design and Implementation (PLDI16). DOI: 10.1145/2908080.2908094

# Science Highlights (Continued)

## Indian summer monsoon response to a range of black carbon aerosols radiative forcing

*Mahesh Kovilakam and Salil Mahajan*

Assessed the range of Indian Monsoon response within the bounds of present day uncertainty of black carbon aerosols (BC) radiative forcing and beyond - into nuclear winter realms.

While black carbon aerosols (BC) are believed to modulate the Indian monsoons, the radiative forcing estimate of BC suffers from large uncertainties globally. We analyze a suite of idealized experiments forced with a range of BC concentrations that span a large swath of the latest estimates of its global radiative forcing. Within those bounds of uncertainty, summer precipitation over the Indian region increases nearly linearly with the increase in BC burden. The linearity holds even as the BC concentration is increased to levels resembling those hypothesized in nuclear winter scenarios, despite large surface cooling over India and adjoining regions. The enhanced monsoonal circulation is associated with a linear increase in the large-scale meridional tropospheric temperature gradient. The precipitable water over the region also increases linearly with an increase in BC burden, due to increased moisture transport from the Arabian sea to the land areas. The wide range of Indian monsoon response elicited in these experiments emphasizes the need to reduce the uncertainty in BC estimates to accurately quantify their role in modulating the Indian monsoons. The increase in monsoonal circulation in response to large BC concentrations contrasts earlier findings that the Indian summer monsoon may break down following a nuclear war.

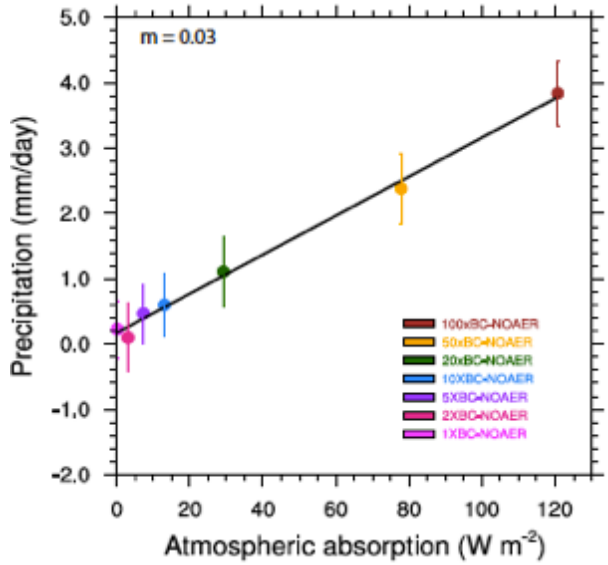


***Linear increase in Indian Monsoon precipitation response to BC forcing***

This work was sponsored by DOE-BER and the Oak Ridge Leadership Computing Facility.

Publication: Kovilakam M. and S. Mahajan (2016): Confronting the 'Indian summer monsoon response to black carbon aerosols' with the uncertainty in its radiative forcing and beyond, Journal of Geophysical Research - Atmospheres (accepted), 121, doi: 10.1002/2016JD024866

# Science Highlights (Continued)

## OpenACC to FPGA: A Framework for Directive-based High-Performance Reconfigurable Computing

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

Researchers proposed an OpenACC-to-FPGA translation framework, which performs source-to-source translation of the input OpenACC program into an output OpenCL code, which is further compiled to an FPGA program by the underlying backend Altera OpenCL compiler.

In this work, we present a directive-based, high-level programming framework for high-performance reconfigurable computing. It takes a standard, portable OpenACC C program as input and generates a hardware configuration file for execution on FPGAs. We implemented this prototype system using our open-source OpenARC compiler; it performs source-to-source translation and optimization of the input OpenACC program into an OpenCL code, which is further compiled into a FPGA program by the backend Altera Offline OpenCL compiler. Internally, the design of OpenARC uses a high-level intermediate representation that separates concerns of program representation from underlying architectures, which facilitates portability of OpenARC. In fact, this design allowed us to create the OpenACC-to-FPGA translation framework with minimal extensions to our existing system. In addition, we show that our proposed FPGA-specific compiler optimizations and novel OpenACC pragma extensions assist the compiler in generating more efficient FPGA hardware configuration files. Our empirical evaluation on an Altera Stratix V FPGA with eight OpenACC benchmarks demonstrates the benefits of our strategy. To demonstrate the portability of OpenARC, we show results for the same benchmarks executing on other heterogeneous platforms, including NVIDIA GPUs, AMD GPUs, and Intel Xeon Phis. This initial evidence helps support the goal of using a directive-based, high-level programming strategy for performance portability across heterogeneous HPC architectures.

Altera/Nallatech provided the FPGA board as an evaluation testbed. This material is based upon work supported by the U.S. Department of Energy (DOE), Office of Science,

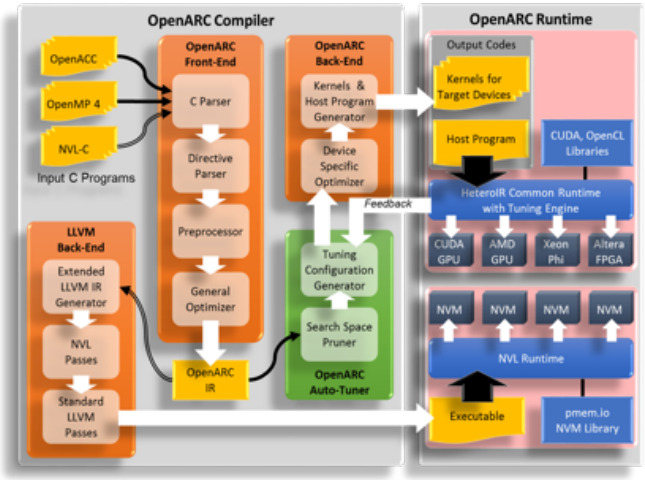


***OpenARC is an open-sourced, very High-level Intermediate Representation (HIR)-based, extensible compiler framework, where various performance optimizations, traceability mechanisms, fault tolerance techniques, etc., can be built for better debuggability/performance/resilience on the complex accelerator computing.***

***Available at http://ft.ornl.gov/research/openarc***

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

Publication: "OpenACC to FPGA: A Framework for Directive-based High-Performance Reconfigurable Computing"; Seyong Lee, Jungwon Kim, and Jeffrey S. Vetter; IEEE International Parallel and Distributed Processing Symposium (IPDPS). DOI:10.1109/IPDPS.2016.28, 2016

# Science Highlights (Continued)

## Paring in a Dry Fermi Sea

*T.A. Maier, P. Staar, V. Mishra, U. Chatterjee, J.C. Campuzano and D.J. Scalapino*

Quantum Monte Carlo simulations of a Hubbard model and an analysis of experimental photoemission data reveal that superconductivity arises in the pseudogap regime from an increase in the strength of the spin-fluctuation pairing interaction as the temperature is lowered.

In the traditional Bardeen–Cooper–Schrieffer theory of superconductivity, the amplitude for the propagation of a pair of electrons with momentum k and -k has a log singularity as the temperature decreases. This so-called Cooper instability arises from the presence of an electron Fermi sea. It means that an attractive interaction, no matter how weak, will eventually lead to a pairing instability. However, in the pseudogap regime of the cuprate superconductors, where parts of the Fermi surface are destroyed, this log singularity is suppressed, raising the question of how pairing occurs in the absence of a Fermi sea. This work used numerical quantum Monte Carlo simulations of a Hubbard model with parameters relevant for the cuprates and an analysis of angular-resolved photoemission experiments on a cuprate superconductor to study this question. In agreement with experiments, the calculations find that the pseudogap suppresses the usual BCS logarithmic instability. In contrast to the traditional theory, however, there is an increase in the strength of the spin–fluctuation pairing interaction as the temperature is lowered, and this leads to a pairing instability. This theoretical prediction is confirmed by the analysis of the experimental photoemission data for two different temperatures on the same system. This demonstrates that in spite of the suppression of the usual BCS logarithmic instability by the pseudogap, the increase in strength of the spin-fluctuation interaction as the temperature is reduced is sufficient to lead to superconductivity.

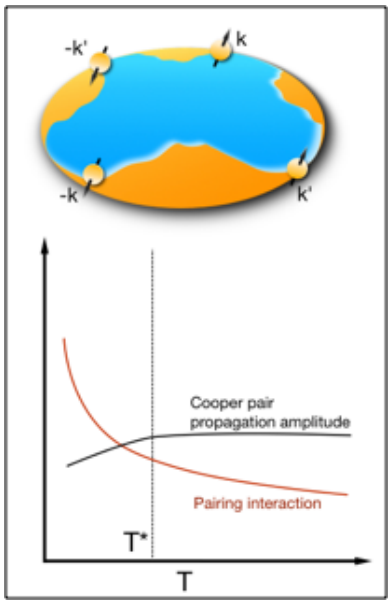


***When the Fermi sea is partially "dried out" in the pseudogap state of the cuprates, the Cooper log divergence is suppressed. In this case, pairing arises from an increase in the strength of the spin-fluctuation pairing interaction as the temperature decreases.***

Work was performed at the Center for Nanophase Materials Sciences, ORNL. A portion of this research was conducted at the Center for Nanophase Materials Sciences, which is sponsored at Oak Ridge National Laboratory by the Division of Scientific User Facilities, U.S. Department of Energy. Part of this research was 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. An award of computer time was provided by the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) programm. This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract DE-AC05-00OR22725.

Publication: T.A. Maier, P. Staar, V. Mishra, U. Chatterjee, J.C. Campuzano and D.J. Scalapino, "Pairing in a dry Fermi sea", Nat. Comm., in print (2016).

# Science Highlights (Continued)

## Greening of the Earth and Responsible Driving Mechanisms

*Jiafu Mao*

Researchers developed an understanding of how dynamics of terrestrial vegetation are responding to global environmental change.

Global environmental change is rapidly altering the dynamics of terrestrial vegetation, with consequences for the functioning of the Earth system and provision of ecosystem services. Yet how global vegetation is responding to the changing environment is not well established. Here we use three long-term satellite leaf area index (LAI) records and ten global ecosystem models to investigate four key drivers of LAI trends during 1982–2009. We show a persistent and widespread increase of growing season integrated LAI (greening) over 25% to 50% of the global vegetated area, whereas less than 4% of the globe shows decreasing LAI (browning). Factorial simulations with multiple global ecosystem models suggest that $CO_2$ fertilization effects explain 70% of the observed greening trend, followed by nitrogen deposition (9%), climate change (8%) and land cover change (LCC) (4%). $CO_2$ fertilization effects explain most of the greening trends in the tropics, whereas climate change resulted in greening of the high latitudes and the Tibetan Plateau. LCC contributed most to the regional greening observed in southeast China and the eastern United States. The regional effects of unexplained factors suggest that the next generation of ecosystem models will need to explore the impacts of forest demography, differences in regional management intensities for cropland and pastures, and other emerging productivity constraints such as phosphorus availability.

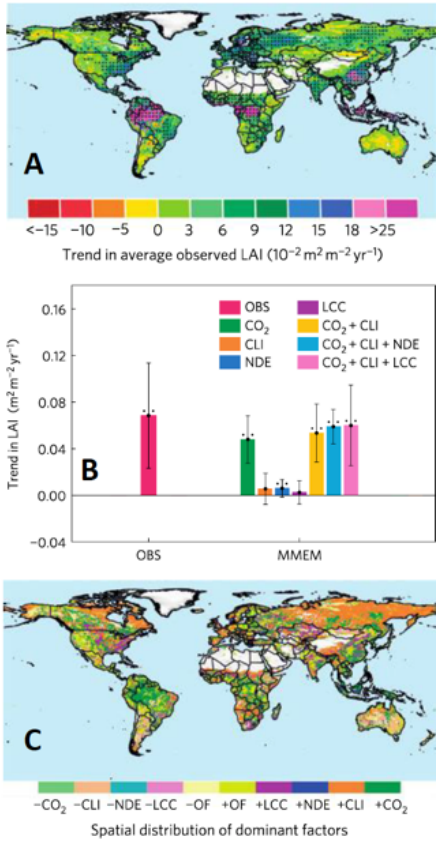


J. Mao is supported by DOE Office of Science, Biological and Environmental Research, including support from the following programs: Terrestrial Ecosystem Science Program (ORNL TES SFA) and Regional and Global Climate Modeling Program (ORNL BGC-Feedbacks SFA)

Publication: Zhu, Z., S. Piao, R. B. Myneni, M. Huang, Z. Zeng, J. G. Canadell, P. Ciais, S. Sitch, P. Friedlingstein, A. Arneth, C. Cao, L. Cheng, E. Kato, C. D. Koven, Y. Li, X. Lian, Y. Liu, R. Liu, J. Mao, Y. Pan, S. Peng, J. Peñuelas, B. Poulter, T. A. M. Pugh, B. D. Stocker, N. Viovy, X. Wang, Y.-P. Wang, Z. Xiao, H. Yang, S. Zaehle, and N. Zeng (2016), Greening of the Earth and its drivers, Nature Clim. Change, doi:10.1038/nclimate3004.

# Science Highlights (Continued)

## Finding Solidifies Optimism for Flexible Electronics

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

Researchers identified an atomic-scale electronic variation related to deformation and edge reconstruction of cleaved black phosphorus, a promising electrical semiconductor.

An increasing number of alternative two-dimensional (2D) materials are being explored in the " post-graphene age," such as transition metal dichalcogenides (TMDs), silicene, and germanene. Among these 2D layered materials, layered black phosphorus (phosphorene) is expected to exhibit superior mechanical, electrical, and optical properties due to an intrinsic and tunable bandgap. It is, besides graphene, the only stable elemental 2D material that can be mechanically exfoliated. Its high hole mobility and direct semiconducting bandgap fuel hope for the development of new electronic devices in the post-silicon era. However, before phosphorene can be integrated in everyday electronics, many aspects of the material remain to be elucidated. For example, the direct bandgap of single layer black phosphorus predicted theoretically, has not been directly measured in experiments, due in part to the difficulty of isolating a single layer of phosphorene. Furthermore, the properties of its edges, which are likely different from those deep in the basal plane, have not been considered in detail. Here, we present a extensive scanning tunneling microscopy/spectroscopy (STM/S) investigation, corroborated by ab initio calculations, to reveal the presence of a semiconducting 2 eV bandgap observed by STS, which represents the direct bandgap predicted for an isolated surface layer of the cleaved solid. In addition, we study a mono-step of layered black phosphorus and demonstrate that the step introduces nontrivial edge states, exhibiting multiple peaks in STS, which can be explained as a signature of one-dimensional (1D) van Hove singularities resulting the reconstructed atoms at the edge (e.g., P atoms at the edge reorder to become coordinated with 3 neighboring P atoms).

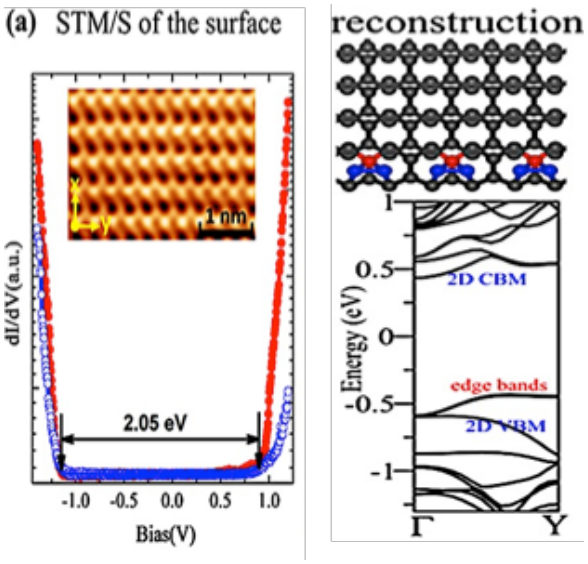


***(a) Scanning Tunneling Spectroscopy shows the energy gap in black phosphorus (BP). Red and blue dotted curves show change in current divided by change in voltage at different surface locations. The insert is a high resolution Scanning Tunneling Microscopy image of BP. (b) Calculations showing atomic reconstruction at a zigzag edge where P atoms (red) adsorb two dangling edge atoms (blue) to have five neighbors, with all edge atoms being three-coordinated (top). Upon edge reconstruction, the edge bands that were originally metallic become semiconducting (bottom frame).***

Work was performed at the Center for Nanophase Materials Sciences at ORNL and at Rensselaer Polytechnic Institute.

Publication: "Electronic Bandgap and Edge Reconstruction in Phosphorene Materials" Liangbo Liang, Jun Wang, Wenzhi Lin, Bobby G. Sumpter, Vincent Meunier, and Minghu Pan, Nano Lett. 2014, DOI: 10.1021/nl502892t

# 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 Kilbey, II*

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

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.

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 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.

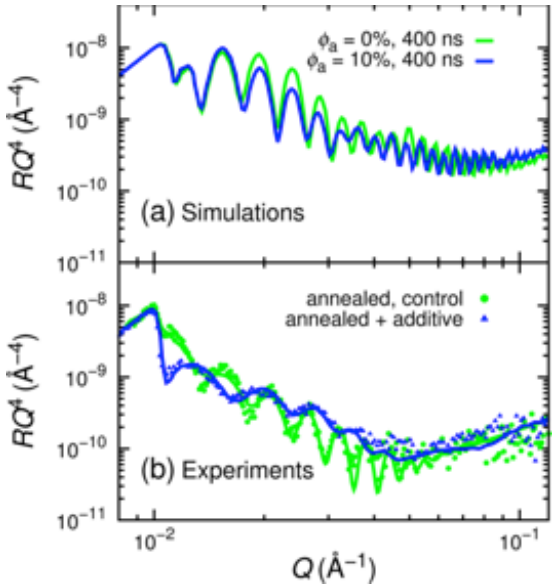


***(Top) Coarse-grained simulation of OPV interface compared (bottom) with neutron reflectivity measurements.***

Publication: 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 Kilbey, II, Petascale Simulations of the Morphology and the Molecular Interface of Bulk Heterojunctions, ACS Nano, DOI: acsnano.6b03009 (2016)

# Science Highlights (Continued)

## Understanding Structure-function Evolution in Complex Solutions of Polymers

*Jim Browning, Ilia N. Ivanov, J. Zhu, N. Herath, K. Hong, Valeria Lauter, Rajeev. Kumar, Bobby Sumpter, Hassina Bilheux, Jim Browning, Changwoo Do, Benjamin Doughty, Yingzhong Ma*

Researchers developed and utilized an unique environmental chamber for in-situ multimodal interrogation with direct feedback to data analytics and advanced simulations that enabled achieving a new level of control of polymer/small molecule assembly in solution and thin films.

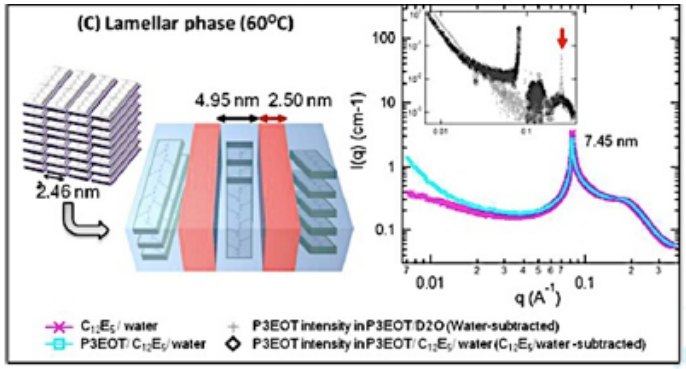


Researchers developed a sample environment where neutron and optical probes worked in concert to track molecular dynamics in solution and in thin films. They used their new capability—in situ multimodal interrogation of structures and functions—to make discoveries likely to improve photovoltaic materials, including assembling a polymer into functional structures in aqueous solution using a detergent-like molecule as a template and improving a solid's performance via an additive. Once placed in the test chamber, a sample can be transported to different instruments for measurements. The chamber has a transparent face to allow entry of laser beams to probe materials. Probing modes include neutrons, X-rays, photons, electrical charge, magnetic spin, and calculations aided by high-performance computing; they can operate simultaneously to characterize matter under a broad range of conditions. The chamber can track molecules under changing temperature, pressure, humidity, light, solvent composition, etc. It extends the probing of materials from typical environments under resting conditions to operating environments that allow scientists to assess how working materials change over time and explore ways to improve their performance.

This work was sponsored by the Center for Nanophase Materials Sciences (CNMS)

Publication: Scientific Reports 5: 13407 (2015), Nanoscale 7, 15007 (2015)

## Adding Fault Tolerance to NPB Benchmarks Using ULFM

*Z. Parchman, G. Vallee, T. Naughton, Christian Engelmann, David Bernholdt, Stephen Scott*

We modified some of the benchmarks of the NAS parallel benchmarks (NPB) to include support of the MPI fault tolerant User Level-Fault Mitigation capabilities as well as application-level strategies and mechanisms for application-level failure recovery. As such, we contributed: (i) an application-level library to "checkpoint" and restore data, (ii) extensions of NPB benchmarks for fault tolerance based on different strategies, (iii) a fault injection tool, and (iv) some preliminary results that show the impact of such fault tolerant strategies on the application execution.

In the world of high-performance computing, fault tolerance application resilience are becoming some of the primary concerns because of increasing hardware failures and corruptions. While the research community has been investigating various options, from system-level solutions to application-level solutions, standards such as the Message Passing Interface (MPI) are also starting at including such capabilities. The current proposal for MPI fault tolerant is based on the User-Level Failure Mitigation (ULFM) concept, which provides means for fault detection and recovery of the MPI layer. This approach does not address application-level recovery, which is current left to application developers. In this work, we present a modification of some of the benchmarks of the NAS parallel benchmark (NPB) to include support of the ULFM as well as application- level strategies and mechanisms for application-level failure recovery. As such, we present: (i) an application-level library to ``checkpoint'' data, (ii) extensions of NPB benchmarks for fault tolerance based on different strategies, (iii) a fault injection tool, and (iv) some preliminary experiments that shows the impact of such fault tolerant strategies on the application execution.

This research was sponsored by and used OLCF resources.

Publication: Proceedings of 25th ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC) 2016: 6th Workshop on Fault Tolerance for HPC at eXtreme Scale (FTXS) 2016

# Science Highlights (Continued)

## Untangling the Effects of Chain Rigidity on the Structure and Dynamics of Strongly Adsorbed Polymer Melts

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

We demonstrate, through computer simulations, that the range of the interfacial layer increases as the stiffness of the polymer chain increases.

A detailed analysis of coarse-grained molecular dynamics simulations of semiflexible polymer melts in contact with a strongly adsorbing substrate is used to characterized the segments in the interfacial layer. For more rigid chains, a tail and an adsorbed segment (a train) dominate while loops are more prevalent in more flexible chains. For segmental dynamics, our results show that it is slower for stiffer chains, and there is a strong correlation between the structure and dynamics in the interfacial layer. However, there is no "glassy layer", and the slowing down in dynamics of stiffer chains in the adsorbed region can be attributed to the densification and a more persistent layering of segments.

Polymers present a broad class of materials that are increasingly utilized in energy generation and conversion, energy storage and lightweight materials, chemical separations and gas purification. Yet a fundamental understanding of the interactions that control macroscopic properties remains limited, especially in multicomponent systems. Building on prior progress on the equilibrium structural aspects of these materials, the overarching goal of the proposed research is to develop a fundamental and predictive understanding of dynamics in multicomponent polymeric materials. The pervasive presence of interfaces between different phases or species in multicomponent systems, and their effect on bulk properties, is the unifying aspect of our research. We address the following fundamental issues: How do the nanoparticle-polymer interactions, polymer rigidity and particle softness affect the structure and dynamics in the interfacial region? How does spatial confinement and hard and soft interfaces modify bulk polymer and nanoparticle diffusion and collective viscoelasticity? How do molecular architecture and morphology of segregated block copolymers affect structure and dynamics of the interfacial region and macroscopic viscoelastic properties of these materials? How can one understand, control and exploit the

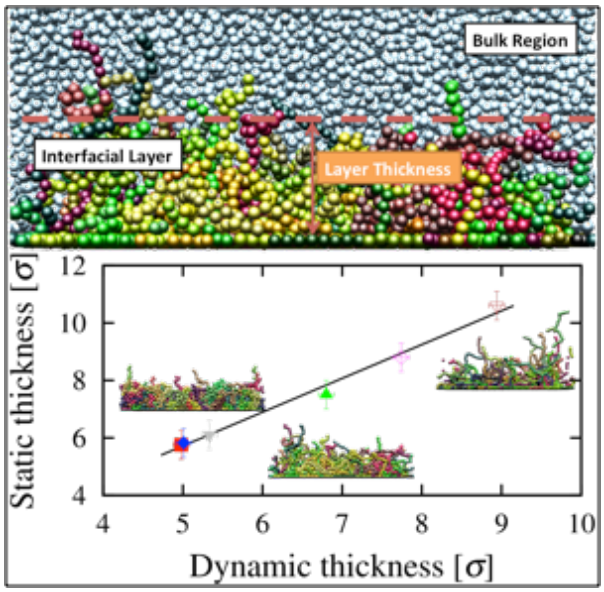


***The interfacial layer is defined as the region where the properties of the chain differs from the bulk (top image). The layer thickness increases as chain rigidity increases. There is direct correlation in the static and dynamics layer thicknesses and this can be explained by the densification of the system in the vicinity of the adsorbed layer, which promotes the slowing down of the polymers (bottom image).***

often pervasive presence of non-equilibrium effects in these macromolecular materials? We will pursue a comprehensive interdisciplinary approach led by advanced theory and simulations, precise synthesis, and state-of-the-art characterization with special emphasis on neutron scattering. The fundamental knowledge developed in this program will contribute to the scientific foundation for the rational design of multicomponent polymer based materials with superior properties and function that will have broad applications in energy technologies.

This research was sponsored by the Department of Energy, Basic Energy Sciences, Materials Science and Engineering Division (MSED); the Center for Nanophase Materials Sciences (CNMS); and the Oak Ridge Leadership Computing Facility (OLCF)

Publication: "Untangling the Effects of Chain Rigidity on the Structure and Dynamics of Strongly Adsorbed Polymer Melts"; : Macromolecules 2015, 48, 4207–4219

# Science Highlights (Continued)

## A Three-dimensional Meso-macroscopic Model for Li-Ion Intercalation Batteries

*S. Allu, S. Kalnaus, S. Simunovic, J. Nanda, J. A. Turner, and S. Pannala*

Researchers performed largest coupled electrochemistry-thermal simulation of a Li-ion battery pack.

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.

The Vehicle Technologies Program under the Office of Energy Efficiency and Renewable Energy has sponsored this research at Oak Ridge National Laboratory. 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 non-exclusive, paid-up, irrevocable, world-wide 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).

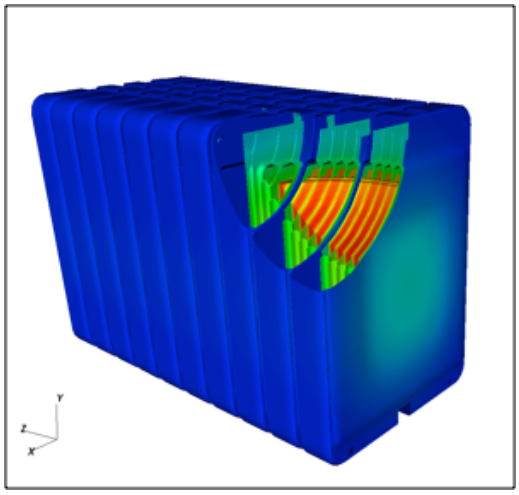

***Pack with 10 modules, each module consisting of 4 cells: Each cell has 32 cell sandwiches stitched with current collectors; Total of 16 million degrees of freedom; ~400k DOFs / cell; multiple zones per layer; 1 hr. simulation time; 24 hours on 1024 cores***

Publication: S. Allu, S. Kalnaus, S. Simunovic, J. Nanda, J. A. Turner, and S. Pannala, “A three-dimensional meso-macroscopic model for Li-Ion intercalation batteries,” Journal of Power Sources, vol. 325, pp. 42–50, Sep. 2016. doi:10.1016/j.jpowsour.2016.06.001.

# News

## Increasing battery safety through high performance computing

*John Turner*

Preventing thermal runaway, or the rapid overheating of a battery that can lead to fires or even explosions, is the focus of several efforts led by John Turner of ORNL's Computational Engineering and Energy Sciences Group. Turner and his team are developing advanced numerical tools to simulate electric vehicle batteries in crashes to better understand the triggers for thermal runaway and to evaluate new battery technologies.

Standard tests of battery stability—like the nail penetration test where a metallic nail is driven through a charged cell with any resulting smoke or flame indicating a short or thermal runaway—are not truly representative of real-world conditions. There have been cases where lithium-ion batteries in cars involved in minor fender benders have appeared to be stable only to later develop problems. Said Turner, "They don't have to be crushed to be damaged inside." Clearly a better understanding of how batteries behave in abusive situations, particularly under mechanical abuse like that encountered in vehicle accidents, is needed.

Developing software to virtually crash an automobile and predict the state of the battery afterwards is one of the goals of a new multi-lab effort funded by the Department of Energy (DOE) Vehicle Technologies Office (VTO). The Consortium for Advanced Battery Simulation (CABS), a collaborative effort that includes Sandia National Laboratories and Lawrence Berkeley National Laboratory, is developing an integrated software environment that includes the effects of microstructure and full geometric complexity.

In addition, as part of a multi-year partnership with the National Highway Traffic Safety Administration (NHTSA), Turner and his team are conducting physical experiments with battery pouch cells, crushing and pinching them to varying amounts and peeling them open to observe effects. The experimental data is being used to validate computational models and to gain insights that could inform future regulations about battery evaluation procedures.

A third related project, also funded by VTO, is led by Ford Motor Company and is focused on simulating the effects of mechanical abuse, with emphasis on commercial software tools such as LS-DYNA.

"High performance computing and these collaborations between government, industry, and national labs are accelerating the pace of discoveries related to battery safety and performance," said Turner, "but there are still many questions to answer."

Members of the energy storage modeling and simulation team include Srikanth Allu, Wael Elwasif, Sergiy Kalnaus, Abhishek Kumar, Damien Lebrun-Grandie, Srdjan Simunovic, Bruno Turcksin, and Hsin Wang.

## Climate study finds human fingerprint in Northern Hemisphere greening

Jiafu Mao

A multinational team led by the Department of Energy's Oak Ridge National Laboratory Climate Change Science Institute has found the first positive correlation between human activity and enhanced vegetation growth.

The research team, led by Jiafu Mao of the Ecosystem Simulation Science group in the Environmental Sciences Division, used new environmental data and strict statistical methods to discover a significant human-vegetation interaction in the northern extratropical latitudes, the section of the planet spanning 30 to 75 degrees north, roughly between the Tropic of Cancer and the North Frigid Zone above the Arctic Circle.

http://www.eurekalert.org/pub_releases/2016-06/drnl-csf062916.php

## Science community scrutinizes p-values, reproducibility

*George Ostrouchov*

P-values have been in the news lately, with at least one journal announcing that it would no longer publish papers containing them because the statistics were too often used to support lower-quality research.

The American Statistical Association has released a statement on p-values, an unprecedented action by the ASA in its 177-year history, says the Computer Science & Mathematics Division's George Ostrouchov. The paper is currently an open access accepted manuscript in The American Statistician that may still undergo minor edits.

George says, "From a statistician's perspective, p-values and statistical significance summarize a complex and narrow situation rendering their interpretation to be highly context dependent. Unfortunately, in our quest to simplify things, they are often used to incorrectly justify broader statements, leading to the current reproducibility crisis in many fields. My favorite commentary on the ASA statement is at FiveThirtyEight."

More articles and blogs about the ASA statement can be found at Following up on the ASA Statement on P-Values and Statistical Significance.

ORNL researchers may continue to use P-values.

# Community Service

- Ferrol Aderholdt, program committee, 9th Workshop on Resiliency in High-Performance Computing (Resilience) at the 22nd International European Conference on Parallel and Distributed Computing (Euro-Par), Grenoble, France, August 22-26, 2015
- Rizwan A. Ashraf, peer reviewer, IEEE Transactions on Evolutionary Computation
- David E. Bernholdt, peer reviewer, Journal of Computational Science
- Swen Boehm, program committee, IEEE Cluster 2016, Taipei, Taiwan, September 12--16, 2016
- Michael J. Brim, program committee, IEEE Cluster 2016, Taipei, Taiwan, September 12--16, 2016
- Christian Engelmann, program committee, 6th International Workshop on Runtime and Operating Systems for Supercomputers (ROSS) at 25th International Symposium on High Performance Distributed Computing (HPDC), Kyoto, Japan, May 31--June 4, 2016
- Christian Engelmann, peer reviewer, International Journal of High Performance Computing Applications (IJHPCA)
- Christian Engelmann, peer reviewer, IEEE Transactions on Parallel and Distributed Systems (TPDS)
- Christian Engelmann, peer reviewer, Journal of Parallel and Distributed Computing (JPDC)
- Christian Engelmann, program committee, 11th International Conference on Availability, Reliability and Security (ARES), Salzburg, Austria, August 31--September 2, 2016
- Christian Engelmann, program committee area chair for System Software, 29th ACM/IEEE International Conference on High Performance Computing, Networking, Storage and Analysis (SC), Salt Lake City, UT, USA, November 13-18, 2016
- Christian Engelmann, program committee co-chair, 9th Workshop on Resiliency in High-Performance Computing (Resilience) at the 22nd International European Conference on Parallel and Distributed Computing (Euro-Par), Grenoble, France, August 22-26, 2015
- Christian Engelmann, program committee, 23rd European MPI Users Group Meeting (EuroMPI), Edinburgh, UK, September 26-28, 2016
- Oscar Hernandez, program committee chair, International Workshop on OpenPOWER for HPC (IWOPH16), at ISC High Performance, Frankfurt, Germany, 23 June 2016
- Oscar Hernandez, steering committee, First International Workshop on Performance Portable Programming Models for Accelerators (P^3MA), at ISC High Performance, Frankfurt, Germany, 23 June 2016.
- Oscar Hernandez, program committee, PGAS Applications Workshop, at the 29th ACM/IEEE International Conference on High Performance Computing, Networking, Storage and Analysis (SC), Salt Lake City, UT, USA, November 13-18, 2016
- Oscar Hernandez, steering committee , Third Workshop on Accelerator Programming using Directives (WACCPD) at the 29th ACM/IEEE International Conference on High Performance Computing, Networking, Storage and Analysis (SC), Salt Lake City, UT, USA, November 13-18,
- 2016
- Terry Jones was invited to lead a chapter of the Time Synchronization Task Force for the North American SynchroPhasor Initative (NASPI) in June of 2016.
- Tonglin Li, peer reviewer, Concurrency and Computing: Practice and Experience
- Tonglin Li, peer reviewer, Journal of Multimedia Tools and Applications (MTAP)
- Tonglin Li, peer reviewer, Frontiers of Information Technology & Electronic Engineering, Special Issue on Distributed Computing and Artificial Intelligence
- Tonglin Li, program committee, 2nd International Conference on Cloud Computing Technologies and Applications (CloudTech'16), Marrakesh, Morocco, May 24--26 2016
- Tonglin Li, program committee, 6th IEEE International Symposium on Cloud and Service Computing (SC2' 16), Nadi, Fiji, December 8--10, 2016
- Swaroop Pophale, program committee, International Workshop on OpenPOWER for HPC (IWOPH'16) at ISC High Performance, Frankfurt, Germany, 23 June 2016
- Swaroop Pophale, program committee chair, OpenSHMEM 2016: Third workshop on OpenSHMEM and Related Technologies, Baltimore, Maryland, August 2-4,2016
- Gregory Watson, program committee, IEEE Cluster 2016, Taipei, Taiwan, September 12--16, 2016
- Jeffrey Vetter, Appointed Associate Editor of ACM Journal on Emerging Technologies in Computing (JETC)
- Jeffrey Vetter, Steering Committee, International Supercomputing Conference (ISC), 2016 - present.
- Jeffrey Vetter, Co-chair, First International Workshop on Post Moores Supercomputing, 2016.
- Jeffrey Vetter, Program Committee, Technical Papers, SC16, 2016.
- Jeffrey Vetter, Program Committee, ICS, 2016.
- Jeffrey Vetter, Program Committee, ISC, 2016.
- Jeffrey Vetter, Organizing Committee, PMMA16, 2016.
- Jeffrey Vetter, Organizing Committee, MODSIM, 2016.

# Community Service (Continued)

- Jeffrey Vetter, Program Committee, High Performance Extreme Computing Conference, 2016.
- Seyong Lee, External reviewer for the International Journal of High Performance Computing
- Jungwon Kim, Referee for ACM Transactions on Architecture and Code Optimization (TACO)
- Jungwon Kim, Referee for Journal of Parallel and Distributed Computing (JPDC)
- M. Graham Lopez, Program Committee Co-chair, International Workshop for OpenPOWER in HPC (IWOPH), held in conjunction with ISC'16, Frankfurt, Germany, June 2016
- M. Graham Lopez, committee member, EuroPAR manycore/multicore track
- M. Graham Lopez, SC'16 Tutorials committee member

# Events

### PAST

### P3MA Workshop

*M. Graham Lopez*

First International Workshop on Performance Portable Programming Models for

Accelerators (P^3MA) co-located with ISC 2016 was held at Frankfurt, Germany on June 23. The workshop solicited papers on topics covering feature sets of programming models (including but not limited to directives-based programming models), their implementations, and experiences with their deployment in HPC applications on multiple architectures, performance modeling and evaluation tools, asynchronous task and event-driven execution/scheduling. We received 13 submissions in total. All submitted manuscripts were peer reviewed. The review process was not double blind, i.e., authors were known to reviewers.

Submissions were judged on correctness, originality, technical strength, significance, quality of presentation, and interest and relevance to the conference scope. We chose 8 papers to be published in the workshop proceedings, Springer Verlag Lecture Notes in Computer Science (LNCS) volumes.

The workshop was held on June 23 at ISC and brought together researchers, vendors, users and developers to brainstorm aspects of heterogeneous computing and its various tools and techniques. Around 50 attendees were present to see Dr. Si Hammond from Sandia National Laboratories, USA, give a keynote on

Balancing Productivity, Portability and Performance - The Challenge for Programming Models at Exascale. All of the 8 accepted papers were presented at the workshop with topics ranging from using low-level to high-level programming models for heterogeneous systems, experiences porting legacy code to accelerators, addressing memory requirements, and creating translations from one standard to the other.

Prof. Haohuan Fu, Deputy Head of the National Supercomputing Center in Wuxi and Associate Professor at the Center for Earth System Science, Tsinghua University, China, gave an invited talk on preparing the Community Atmospheric Model climate application to run and scale on Sunway TaihuLight, announced at ISC'16 as the new number one HPC system on the Top500 list. His talk included their experiences with custom loop transformation tools for code refactoring, using OpenACC to program the heterogeneous architecture of TaihuLight, and extensions to the OpenACC standard that were implemented to enable various optimizations in the application.

NVIDIA generously offered to sponsor the 'Best Paper Award' with NVIDIA's newest PASCAL compute capable card. This award was presented to "Task-Based Cholesky Decomposition on Knights Corner using OpenMP" from UTK. The award was determined by the Technical Program Committee and the Program Chairs from viewpoints of the technical and scientific merits, impact on the science and engineering of the research work and the clarity of presentation of the research contents in the paper.

http -//www.csm.ornl.gov/workshops/p3ma2016/committee.html

### CoDA 2016: Conference on Data Analysis – Exploring Data-Focused Research across the Department of Energy (March 2-4, 2016)

*M. Graham Lopez*

Another installment (the third) of the biennial Conference on Data Analysis (CoDA) 2016 took place in Santa Fe, NM, this March. This unique and "just the right size" conference explores data-focused research across the Department of Energy in a three-day single track format with only invited talks and poster sessions. ORNL participation included CSMD's George Ostrouchov on the organizing committee, invited talks by CSED's Budhu Bhaduri on "Big Data, Geospatial Computing, and My 2 Cents in an Open Data Economy" and by Arjun Shankar on "Data-Driven VERDE: Visualizing Energy Resources Dynamically on Earth," as well as a number of posters.

In addition to the single track technical talks and poster sessions, the conference features a fun banquet with an excellent speaker on a popular topic. This year's speaker was University of Chicago's Rayid Ghani, speaking about "Data Science for Social Good." This follows the 2014 banquet talk by Amanda Cox on "Sketching the News: Data Graphics at the New York Times" and the 2012 banquet talk by Graham Walsh on "Blowing Stuff Up for Science: Explosives Experiments with the MythBusters."

# Events (Continued)

### CAM Seminar Series

- Dr. Moulay Abdellah Chkifa, Oak Ridge National Laboratory - New Hierarchical Non-intrusive Colocation Polynomial Methods in High Dimension
- Dr. Xiu Ye, University of Arkansas - Weak Galerkin Method and Its Applications
- Dr. Andrea L'Afflitto, University of Oklahoma - Feedback Optimal Control, Robust Control, and Differential Games - Theory, Numerical Solutions, and Applications
- Dr. Richard Barnard, Oak Ridge National Laboratory - PDE-constrained optimization involving multiswitching and volumetric dose penalization
- Mr. Hrvoje Gotovac and Mr. Veljko Srzic, University of Split - Adaptive Multiresolution Modeling of Groundwater Flow and Transport Subsurface Processes
- Dr. Ajay Jasra, National University of Singapore - Multilevel Sequential Monte Carlo
- Dr. Andrea Bertozzi, UCLA - Geometric Graph-based Methods for High Dimensional Data
- Dr. Qiwei Sheng, University of Tennessee - An Optimal Cascadic Multigrid Method for the Radiative Transfer Equation
- Dr. Feng Bao, Oak Ridge National Laboratory - Hierarchical Optimization for Neutron Scattering Problems
- Mr. Viktor Reshniak, Middle Tennessee State University - Acceleration of the Multilevel Monte Carlo for Certain Classes of Differential Systems
- Dr. Xingjie Li, University of North Carolina - Charlotte - Coupling of Nonlocal and Local Diffusion Through a Divergence Formulation
- Ms. Xiaochuan Tian, Columbia University - Asymptotically Compatible Schemes for Robust Discretization of Nonlocal Models
- Mr. Yunzhe Tao, Columbia University - Robust a Posteriori Stress Analysis of Nonlocal Models via Nonlocal Gradients

### CSMD Seminars

- Jai Dayal - Middleware for Managing Large Scale In Situ Science Analysis Workflows
- Jean-Phillipe A. Thomas - Modeling and prevention of coarse microstructural features in aerospace alloys
- Matthew Wolf - Managing Complex Workflows at Scale and In Situ
- Eugene Dumitrescu - Error detection for reliable direct process tomography and channel estimation
- Vivek Mishra - Manifestation of Incipient Band Superconductivity on Neutron Resonance
- Michael Milinkovich - Open Source Is Mainstream
- Greg Dreifus - The Way Forward in Additive Manufacturing Tool Path Optimizations
- Siddharta Santra - Towards Practical Quantum Computing - Lessons from Quantum Annealing on Dwave
- Lingfei (Teddy) Wu - Algorithms and Software for Large Scale Problems in SVD Computations and in Big Data Applications

## FUTURE

### The 4th Workshop on Sparse Grids and Applications - Miami, Florida, October 4-7, 2016.

Following the highly successful workshops on Sparse Grids and Applications: SGA 2011 in Bonn, Germany, SGA 2012 in Munich, Germany, and SGA2014 in Stuttgart, Germany, the 4th workshop on Sparse Grid and Applications will take place October 4-7, 2016 in Miami, Florida, USA. The workshop is cosponsored by Oak Ridge National Laboratory, and the Office of Science, Advanced Simulation Research Computing (ASCR), at the Department of Energy.

The aim is to bring together researchers and practitioners using sparse grids and their variants, and both theoretical and practical talks are highly welcome.

A detailed description of the workshop can be found at: http://www.csm.ornl.gov/workshops/SGA2016/index.html

### Advancing X-cutting Ideas for Computational Climate Science (AXICCS 2016)

*Kate Evans*

A workshop to discuss bold, new computational ideas to address longer-term science needs for climate modeling is being convened on September 12-13, 2016 in Rockville, MD. As the climate science community explores a host of critical science questions to be answered over the next 10+ years, with the aim of informing stakeholders about the ongoing changes in global and local climate, there is a growing recognition of the expanding requirements for multiscale, global, coupled Earth system models. We expect them to provide much more detail and fidelity, with a much better understanding of their uncertainties, while still executing robustly and efficiently on ever larger and more complex computing systems.

For more information, please go to www.csm.ornl.gov/workshops/AXICCS

### 4th Workshop on Sparse Grids and Applications

*Clayton Webster*

Sparse grids are a popular approach for the numerical treatment of high-dimensional problems. Where classical numerical discretization schemes fail in more than three or four dimensions, sparse grids, in their different flavors, are

# Events (Continued)

frequently the method of choice, be it spatially adaptive in the hierarchical basis or via the dimensionally adaptive combination technique.

Following the highly successful workshops on Sparse Grids and Applications: SGA 2011 in Bonn, Germany, SGA 2012 in Munich, Germany, and SGA2014 in Stuttgart, Germany, the 4th workshop on Sparse Grid and Applications will take place October 4-7, 2016 in Miami, Florida. The workshop is co¬-sponsored by the Oak Ridge National Laboratory, and The Office of Science, Advanced Simulation Research Computing (ASCR), at the Department of Energy.

More info here – www.csm.ornl.gov/workshops/SGA2016

### CCSI EARTH SYSTEM MODELING WORKSHOP - Memorial to David Erickson (Earth System Modeler – ORNL)

*Kate Evans*

The Computational Earth Sciences group within the Computer Science and Mathematics Division and the Earth System Modeling theme (ESM) theme within the Climate Change Science Institute is conducting a day-long workshop in honor of their treasured colleague David Erickson, who passed away on November 16, 2015. The theme of the workshop is to present and discuss research in the area that Dave contributed most significantly.

The goals of this year's workshop are to educate participants about global Earth system model development and analysis and generate new ideas and collaborations, both of which were priorities for Dave. The outcome will be that visiting collaborators of Dave and their students and CCSI and other ORNL researchers can grab the baton of Dave's research area and foster additional high profile climate science research. We aim to engage those interested in joining our efforts from other areas of expertise at ORNL (e.g. ASCR), so they can best advise us in future work and identify potential collaborations.

More info here – www.csm.ornl.gov/workshops/esm2016

# Software

### The ILAMB Benchmarking System (v2) Release

*Nathaniel Collier and Forrest M. Hoffman (ORNL), Mingquan Mu and James T. Randerson (University of California at Irving), William J. Riley and Charles D. Koven (LBNL), Gretchen Keppel-Aleks (University of Michigan), and David M. Lawrence (National Center for Atmospheric Research).*

The SFA has collected a large number of observational data (60 datasets) across a wide swath of measurable quantities from land models (24 variables). The release of this software puts this information as well as the collective benchmarking expertise of the SFA team in the hands of the community. In addition to the direct dissemination of our research, the software framework itself is generic and extensible, which allows the community to integrate datasets and comparisons that are relevant to their own projects. The ease in which this may be done positions the ILAMB software package to be part of model development workflow in modeling centers across the globe.

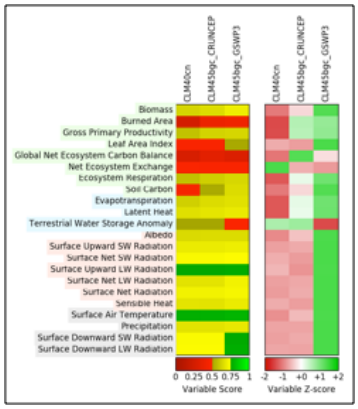


***Summary graphic generated by the ILAMB package depicting model performance across a wide variety of variables, emphasizing absolute performance (left) as well as relative performance (right)***

This research was performed for the Biogeochemistry--Climate Feedbacks Scientific Focus Area, which is sponsored by the Regional and Global Climate Modeling (RGCM) Program in the Climate and Environmental Sciences Division (CESD) of the Biological and Environmental Research (BER) Program in the U.S. Department of Energy Office of Science.

# Awards

### Early Career - Travis Humble and Clayton Webster

The Early Career Research Program, now in its eighth year, supports the development of individual research programs of outstanding scientists early in their careers and stimulates research careers in the disciplines supported by the DOE Office of Science. Opportunities exist in the following program areas: Advanced Scientific Computing Research (ASCR); Biological and Environmental Research (BER); Basic Energy Sciences (BES), Fusion Energy Sciences (FES); High Energy Physics (HEP), and Nuclear Physics (NP).

***Dr. Travis Humble***

For the project "Accelerating Applications of High-Performance Computing with Quantum Processing Units."

Quantum computing promises new approaches for solving hard computations by using the quantum physical processes found in atoms and molecules. Dr. Humble's research will investigate how quantum computing platforms can be leveraged to support scientific computing at DOE High Performance Computing (HPC) facilities. The research will assess the potential for quantum computing to accelerate scientific applications in computational chemistry, materials science, and data analytics. The project will translate scientific software into a representation that can run on computer systems that host both conventional and quantum processing units. This model of hybrid computation requires the development of novel run-time and system execution models to manage and process quantum programs. Developing expectations for future quantum processing units allows HPC stakeholders to evaluate the merits of quantum computing when planning next-generation systems.

***Dr. Clayton Webster***

For the project "Mathematical Methods for Optimal Polynomial Recovery of High-Dimensional Systems from Noisy Data."

Dr. Webster will focus on developing polynomial approximation methods for data from physical experiments and numerical simulations. The need for efficient approximations arises in many energy and materials science applications, where constructing the high-dimensional solution map requires repeated measurements from time-consuming experiments or an ensemble of high-dimensional, parameterized numerical simulations. The research pursues advances in constructive approximation theory, convex and constrained optimization, sparse models, data

reduction, and high-dimensional sampling strategies. Innovative methods that exploit sparsity and compression will be developed for inexpensive and accurate approximations that can be used in understanding and analyzing problems arising in plasma physics, molecular electronic

structures, turbulent flows, and other DOE-mission applications.

# Publications of Note

### A Survey of Techniques for Designing and Managing CPU Register File

*Sparsh Mittal, Concurrency and Computation: Practice and Experience, 2016*

Abstract: Processor register file (RF) is an important microarchitectural component used for storing operands and results of instructions. The design and operation of RF has crucial impact on the performance, energy efficiency and reliability of the processor and hence, several techniques have been recently proposed to manage RF in modern processors. In this paper, we present a survey of techniques for architecting and managing CPU register file. We classify the techniques across several parameters to underscore their similarities and differences. We hope that this paper will provide insights to researchers into working of RF and inspire even more efforts towards optimization of RF in next-generation computing systems.

### A Survey Of Cache Bypassing Techniques

*Sparsh Mittal, Journal of Low Power Electronics and Applications, vol. 6, no. 2, pages 5:1-5:30, 2016. http://dx.doi.org/10.3390/jlpea6020005*

Abstract: With increasing core-count, the cache demand of modern processors has also increased. However, due to strict area/power budgets and presence of poor data-locality workloads, blindly scaling cache capacity is both infeasible and ineffective. Cache bypassing is a promising technique to increase effective cache capacity without incurring power/area costs of a larger sized cache. However, injudicious use of cache bypassing can lead to bandwidth congestion and increased miss-rate and hence, intelligent techniques are required to harness its full

# Publications of Note (Continued)

potential. This paper presents a survey of cache bypassing techniques for CPUs, GPUs and CPU-GPU heterogeneous systems, and for caches designed with SRAM, non-volatile memory (NVM) and die-stacked DRAM. By classifying the techniques based on key parameters, it underscores their differences and similarities. We hope that this paper will provide insights into cache bypassing techniques and associated tradeoffs and will be useful for computer architects, system designers and other researchers.

### OpenACC to FPGA: A Framework for Directive-based High-Performance Reconfigurable Computing

*Seyong Lee, Jungwon Kim, and Jeffrey S. Vetter, IEEE International Parallel and Distributed Processing Symposium (IPDPS), 2016. DOI: 10.1109/IPDPS.2016.28*

Abstract: This paper presents a directive-based, high-level programming framework for high-performance reconfigurable computing. It takes a standard, portable OpenACC C program as input and generates a hardware configuration file for execution on FPGAs.

We implemented this prototype system using our open-source OpenARC compiler; it performs source-to-source translation and optimization of the input OpenACC program into an OpenCL code, which is further compiled into a FPGA program by the backend Altera Offline OpenCL compiler.

Internally, the design of OpenARC uses a high-level intermediate representation that separates concerns of program representation from underlying architectures, which facilitates portability of OpenARC. In fact, this design allowed us to create the OpenACC-to-FPGA translation framework with minimal extensions to our existing system.

In addition, we show that our proposed FPGA-specific compiler optimizations and novel OpenACC pragma extensions assist the compiler in generating more efficient FPGA hardware configuration files.

Our empirical evaluation on an Altera Stratix V FPGA with eight OpenACC benchmarks demonstrate the benefits of our strategy.

To demonstrate the portability of OpenARC, we show results for the same benchmarks executing on other heterogeneous platforms, including NVIDIA GPUs, AMD GPUs, and Intel Xeon Phis. This initial evidence helps support the goal of using a directive-based, high-level programming strategy for performance portability across heterogeneous HPC architectures.

### IMPACC: A Tightly Integrated MPI+OpenACC Framework Exploiting Shared Memory Parallelism

*Jungwon Kim, Seyong Lee, and Jeffrey S. Vetter, In Proceedings of the 25th ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC), pages 189-201, 2016. http://dx.doi.org/10.1145/2907294.2907302*

Abstract: We propose IMPACC, an MPI+OpenACC framework for heterogeneous accelerator clusters. IMPACC tightly integrates MPI and OpenACC, while exploiting the shared memory parallelism in the target system. IMPACC dynamically adapts the input MPI+OpenACC applications on the target heterogeneous accelerator clusters to fully exploit target system-specific features. IMPACC provides the programmers with the unified virtual address space, automatic NUMA-friendly task-device mapping, efficient integrated communication routines, seamless streamlining of asynchronous executions, and transparent memory sharing. We have implemented IMPACC and evaluated its performance using three heterogeneous accelerator systems, including Titan supercomputer. Results show that IMPACC can achieve easier programming, higher performance, and better scalability than the current MPI+OpenACC model.

### A Distributed OpenCL Framework using Redundant Computation and Data Replication

*Junghyun Kim, Gangwon Jo, Jaehoon Jung, Jungwon Kim, and Jaejin Lee, In Proceedings of the 37th ACM SIGPLAN conference on Programming Language Design and Implementation (PLDI), pages 553-569, 2016. http://dx.doi.org/10.1145/2908080.2908094*

Abstract: Applications written solely in OpenCL or CUDA cannot execute on a cluster as a whole. Most previous approaches that extend these programming models to clusters are based on a common idea: designating a centralized host node and coordinating the other nodes with the host for computation. However, the centralized host node is a serious performance bottleneck when the number of nodes is large. In this paper, we propose a scalable and distributed OpenCL framework called SnuCL-D for large-scale clusters. SnuCL-D's remote device virtualization provides an OpenCL application with an illusion that all compute devices in a cluster are confined in a single node. To reduce the amount of control-message and data communication between nodes, SnuCL-D replicates the OpenCL host program execution and data in each node. We also propose a new OpenCL host API function and a queueing optimization technique that significantly reduce the overhead incurred by the previous centralized approaches. To show the effectiveness of SnuCL-D, we evaluate SnuCL-D with a microbenchmark and eleven benchmark applications on a large-scale CPU cluster and a medium-scale GPU cluster.

# Publications of Note (Continued)

### NVL-C: Static Analysis Techniques for Efficient, Correct Programming of Non-Volatile Main Memory Systems

*Joel E. Denny, Seyong Lee, and Jeffrey S. Vetter, In Proceedings of the 25th ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC), 2016*

Abstract: Computer architecture experts expect that non-volatile memory (NVM) hierarchies will play a more significant role in future systems including mobile, enterprise, and HPC architectures. With this expectation in mind, we present NVL-C: a novel programming system that facilitates the efficient and correct programming of NVM main memory systems. The NVL-C programming abstraction extends C with a small set of intuitive language features that target NVM main memory, and can be combined directly with traditional C memory model features for DRAM. We have designed these new features to enable compiler analyses and run-time checks that can improve performance and guard against a number of subtle programming errors, which, when left uncorrected, can corrupt NVM-stored data. Moreover, to enable recovery of data across application or system failures, these NVL-C features include a flexible directive for specifying NVM transactions. So that our implementation might be extended to other compiler front ends and languages, the majority of our compiler analyses are implemented in an extended version of LLVM's intermediate representation (LLVM IR). We evaluate NVL-C on a number of applications to show its flexibility, performance, and correctness.

### Improving DRAM Bandwidth Utilization with MLP-Aware OS Paging

*Rishiraj A. Bheda, Jeffrey S. Vetter, Thomas M. Conte *To be published in the ACM Digital Library as part of Memsys 2016 [International Symposium on Memory Systems]*

Optimal use of available memory bank-level parallelism and channel bandwidth heavily impacts the performance of an application. Research studies have focused on improving bandwidth utilization by employing scheduling policies and request re-ordering techniques at the memory controller. However, potential to extract memory performance by intelligent page allocation that maximizes opportunity for bank-level parallelism and row buffer hits is often overlooked. The actual physical page location in memory has a huge impact on bank conflicts and potential for prioritizing low-latency requests such as row buffer hits. We demonstrate that with more intelligent virtual to physical paging mechanisms it is possible to reduce bank conflicts at the memory and achieve higher bandwidth utilization. Such intelligent paging mechanisms can then form a basis for other request re-ordering techniques to further improve memory performance. In this study we only focus on virtual-to-physical paging techniques and demonstrate 38.4% improvement on DRAM bandwidth utilization with a profile based scheme. We study a wide variety of workloads from varied benchmark suites. We present results for profile based as well as preliminary results for dynamically adaptive paging techniques. Our results demonstrate improved bandwidth utilization with DRAM aware page layouts. Dynamic paging schemes further demonstrate the potential of run-time adaptive techniques in improving bandwidth utilization of increasingly parallel multi-channel main memory systems.

# 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