

Seyong Lee

601 Rockwell Farm Lane,
Farragut, TN 37934
Phone : 765-532-0877
Email : lees2@ornl.gov
Web: <http://ft.ornl.gov/~lees2>

Research Interest

Parallel programming and compile-time/runtime performance optimization on emerging hardware architectures including multi-cores and hardware accelerators
Program analysis and optimizing compiler for high-performance computing
Internet computing /Cloud computing and sharing

Education

8/2004 ~ 5/2011

Purdue University (West Lafayette, Indiana)
Ph.D. in Electrical and Computer Engineering (**GPA 4.0/4.0**)
Advisor: Professor Rudolf Eigenmann

8/2002~5/2004

Purdue University (West Lafayette, Indiana)
Master of Science in Electrical and Computer Engineering (**GPA 3.9/4.0**)
Advisor: Professor Rudolf Eigenmann

3/1995~2/1999

Seoul National University (Seoul, South Korea)
Bachelor of Science in Electrical Engineering (**honors**) (**GPA 3.76/4.30 (3.73/4.0)**)
Advisor: Professor Beom Hee Lee

Research Experience

OpenARC: Open Accelerator Research Compiler (<http://ft.ornl.gov/research/openarc>)

- Develop an open-sourced, high-level intermediate representation-based, extensible compiler framework, which allows full research contexts for directive-based accelerator computing.
 - o Support full features of OpenACC V1.0 (+ subset of V2.0, array reductions, and function calls).
 - o Support CUDA/OpenCL back ends (NVIDIA/AMD GPUs, Intel MICs, and Altera FPGAs).
 - o Equipped with various advanced analysis/transformation passes and built-in tuning tools.

Productive GPU Programming Environment

- Evaluate existing directive-based, high-level GPU programming models to get insights on the current research issues and future directions for the productive GPU programming.
- Existing directive-based GPU programming models (PGI Accelerator, HMPP, R-Stream, OpenACC, and OpenMPC) are evaluated using various benchmarks from diverse application domains.

OpenMP to GPU: Automatic translation and adaptation of OpenMP shared-memory programs onto GPUs.

- Developed the compiler system that translates OpenMP-based shared-memory programs into CUDA-based GPGPU programs and optimizes their performance automatically.
- Created a reference tuning framework, which is able to suggest applicable tuning configurations for a given input OpenMP program, generate CUDA code variants for each tuning configuration, and search the best optimizations for the generated CUDA program automatically.

ATune: Compiler-Driven Adaptive Execution

- Created a tuning system, which adaptively optimizes MPI applications in a distributed system.
- This project is parts of a larger effort that aims at creating a global information sharing system, where resources, such as software applications, computer platforms, and information can be shared, discovered, adapted to local needs.

iShare: Internet-sharing middleware and collaboration

- Developed domain-specific ranking and content search mechanisms for P2P-based Grid environment.
- Developed resource-availability-prediction mechanism for fine-grained cycle sharing system.

MaRCO: MapReduce with Communication Overlap

- Developed efficient communication overlapping mechanisms to increase the performance of Google's MapReduce system.
- Implemented the proposed overlapping mechanism in the Apache Hadoop system.

Work Experience

5/2019 ~ present
1/2014 ~ 4/2019
5/2011 ~ 12/2013

Senior R&D Staff, Future Technology Group, Oak Ridge National Laboratory
R&D Staff, Future Technology Group, Oak Ridge National Laboratory
R&D Associate, Future Technology Group, Oak Ridge National Laboratory
(<http://ft.ornl.gov>)

- Develop high-level programming models for future, heterogeneous computing systems.
- Develop Open Accelerator Research Compiler (OpenARC) as the main architect.

9/2009 ~ 12/2009

Software Engineer (Intern), NEEScomm, Discovery Park, Purdue University
(<https://nees.org>)

- Developed a HUBzero-based cloud computing system for NEES (Network for Earthquake Engineering Simulation).
 - Developed web interfaces in the Joomla Content Management System to communicate with Oracle database and NEES data repository.
 - Configured various applications such as Apache HTTP server, Mailman, SVN, Java EE applications running on a Java Application Server (JBoss).

1/1999 ~ 7/2002

Engineer, R&D Center, Xeline Co., Ltd. (<http://www.xeline.com>)
SAMSUNG & Xeline Powerline Home Automation System building project

Affiliation: Samsung Electronics and Xeline

Term : 10/2001 ~ 7/2002

- Developed Home Automation System using Xeline's PLC modems.

CISCO Systems & Xeline Powerline Network building project (CEAD)

Affiliation: Cisco Systems and Xeline

Term : 4/2001 ~ 8/2001

- Developed PCI based Powerline Communication (PLC) card using Xeline's PLC modem chipset and AMD Ethernet MAC Controller Chipset (AM79C971).
- Built Powerline communication network with CISCO Headend Router and Customer Premise Equipment using PCI based PLC card.

Discrete Multi Tone (DMT) Powerline Communication MODEM Design

Term : 1/2001 ~ 7/2002

- DMT Modem development using XILINX FPGA (VERTEX, VERTEXE), TI DSP (TMS320C670), specifically designing the Digital Interface part of the modem including MII, MDIO, and DI Controller.

Multi-channel Quaternary Frequency Shift Keying (QFSK) Powerline Communication MODEM HW Design & Emulation

Term : 4/2000 ~ 12/2000

- Designed physical layer specification of Multichannel QFSK modem.
- Designed Multi-channel QFSK modem simulator using C++ and MATLAB.
- Designed physical layer of modem chipset using VHDL.
- Performed overall hardware emulation using Xilinx FPGA and Analog Front End board.
- Performed Synthesis/Simulation for ASIC implementation.

Teaching Experience

Teaching Assistant of ECE 461 (Software Engineering)

Electrical and Computer Engineering, Purdue University

Term: 1/2005 ~ 5/2005

- Instructed two lab sessions and held office hours to help students.
- Conducted lab management jobs such as account managing, CVS and other utility environment setup, and etc.

Professional Service

Member of the OpenACC Technical Committee and Test-Suite Committee

Member of OpenMP Language Committee and OpenMP Accelerator Subcommittee

Member of SEED Review Committee, Computer Science and Mathematics Division, Oak Ridge National Laboratory

Member of Science Council, Computer Science and Mathematics Division, Oak Ridge National Laboratory, 6/2017 ~ 5/2020

Member of the NVIDIA PathForward Working Group, Exascale Computing Project PathForward Program, 2018 ~ 2020

Award Committee Member for Computer Science and Mathematics Division Awards, 2018 and 2019

Award Committee Member for 2017 IEEE CS TCHPC Award for Excellence for Early Career Researchers in High Performance Computing, 2017

Science and Innovation Culture Metric Committee, Computing and Computational Science Directorate, Oak Ridge National Laboratory, 2016

Guest Editor for the special issue on “High Performance Reconfigurable computing” in Journal of Algorithms, 2019

Guest Editor for the special issue on “Program Analysis and Optimizing Compilers for High-Performance Computing” in Journal of Electronics, 2021

Program Committee Member

- The IEEE Transactions on Parallel and Distributed Systems (TPDS)’s special section on Innovative R&D toward the Exascale Era: 2021
- The International Conference on Parallel Architectures and Compilation Techniques (PACT) (External Review Committee): 2019 and 2020
- The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC): 2018
- SC Workshop Program Committee, the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC): 2021
- ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS): 2018
- ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP): 2014 and 2020 (External Review Committee)
- IEEE International Parallel & Distributed Processing Symposium (IPDPS): 2017, 2018, 2019, and 2021
- International Conference on Parallel Processing (ICPP): 2013, 2020, and 2021
- IEEE International Conference on High Performance Computing, Data, and Analytics (HiPC): 2019
- ISC High Performance: 2019 and 2020 (Research Poster Committee)
- International European Conference on Parallel and Distributed Computing (Euro-Par): 2017 and 2019
- The IEEE International Conference on Parallel and Distributed Systems (ICPADS): 2013, 2014, 2015, 2016, and 2017
- IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid): 2015, 2016, and 2017 (Posters and Research Demos track)
- IEEE International Performance Computing and Communications Conference (IPCCC):

- 2018, 2019, and 2020
- International Conference on Advanced Engineering Computing and Applications in Science (ADVCOMP): 2017 and 2018
- IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA): 2017 and 2018
- International Symposium on Computing and Networking (CANDAR): 2016
- IEEE International Conference on Computational Science and Engineering (CSE): 2020
- International Workshop on Programming Models, Languages, and Compilers for Manycore and Heterogeneous Architectures (PLC): 2015
- International Workshop on Representative Applications (WRAP): 2015, 2017, and 2018
- International Workshop on Accelerator Programming Using Directives (WACCPD): 2014, 2015, 2016, 2017, 2018, and 2019
- International Workshop on Accelerators and Hybrid Exascale Systems (AsHES): 2016, 2017, 2018, 2019, 2020, and 2021
- International Workshop on Legacy HPC Application Migration (LHAM): 2016, 2017, 2018, and 2020
- International Workshop on Computer Systems and Architectures (CSA): 2019
- International Workshop on High-Level Parallel Programming Models and Supportive Environments (HIPS): 2019, 2020, and 2021
- International Workshop on Legacy Software Refactoring for Performance (REFAC): 2019
- International Workshop on Languages and Compilers for a Parallel Computing (LCPC): 2019

External Reviewer (Journals, Conferences, Workshops, and Research Proposals)

- Journals

- Journal of Parallel and Distributed Computing (JPDC): 2009, 2020, and 2021
- IEEE Transactions on Parallel and Distributed Systems (TPDS): 2014, 2016, 2018, and 2020
- IEEE Micro: 2017 and 2021
- International Journal of High Performance Computing Applications (IJHPCA): 2012, 2015, 2016, 2018, and 2020
- International Journal on Parallel Computing (ParCo): 2013, 2015, 2017, 2018, and 2020
- Journal of Algorithms: 2020
- IEEE Access: 2019
- ACM Transactions on Modeling and Performance Evaluation of Computing Systems (ToMPECS): 2015
- International Journal of Computing Science and Applications (CyS): 2015
- ACM Transactions on Architecture and Code Optimization (ACMTACO): 2013 and 2014
- ACM Transactions on Embedded Computing Systems (TECS): 2020
- International Journal of Software and Systems Modeling (SOSYM): 2011
- Journal of Software: Practice and Experience (SPE): 2010, 2019, and 2021
- Journal of Pure and Applied Mathematics (JPAM): 2017
- Journal of Earth Science (JES): 2017
- International Journal of High Performance Computing and Networking (IJHPCN): 2017
- Journal of Computers (Computers): 2017
- Transactions on Computers (TC): 2017
- International Journal of Parallel Programming (IJPP): 2018
- Future Generation Computer Systems (FGCS): 2018
- Integration, the VLSI Journal (VLSI): 2018
- Open Computer Science (OpenCS): 2018
- SoftwareX: 2018 and 2019
- Scientific Programming: 2019
- Mathematics: 2020
- Symmetry: 2020

- Conferences

- International Conference on Parallel Architectures and Compilation Techniques (PACT): 2010 and 2012
- ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI): 2011
- IEEE International Parallel & Distributed Processing Symposium (IPDPS): 2010 and 2013

- International Conference on Supercomputing (ICS): 2008, 2011, 2013, and 2016
- ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC): 2007 and 2013
- International Symposium on Code Generation and Optimization (CGO): 2013 and 2014
- IEEE International Conference on High Performance Computing (HiPC): 2009 and 2010
- International Conference on Distributed Computing Systems (ICDCS): 2006
- ACM/SPEC International Conference on Performance Engineering (ICPE): 2011
- International Conference on Grid and Pervasive Computing (GPC): 2007 and 2008
- Innovative Parallel Computing Foundations & Applications of GPU, Manycore, and Heterogeneous Systems (INPAR): 2012
- International Conference on Parallel Processing (ICPP): 2019
- **Workshops**
 - International Workshop on Languages and Compilers for a Parallel Computing (LCPC): 2006, 2007, 2011, and 2014
 - International Workshop on OpenMP (IWOMP): 2007, 2009, and 2011
 - International Workshop on Advanced Parallel Processing Techniques (APPT): 2011
 - Workshop on Desktop Grids and Volunteer Computing Systems (PCGrid): 2008
 - Workshop on Exploring Parallelism with Transactional Memory and other Hardware Assisted Methods (EPHAM): 2008 and 2009
- **Research Proposals**
 - The General Research Fund, the Research Grants Council of Hong Kong: 2011
 - Department of Energy (DOE) Office of Science Small Business Innovation Research (SBIR) & Small Business Technology Transfer (STTR) program: 2015
 - Natural Sciences and Engineering Research Council of Canada (NSERC): 2019
 - eTEC2020: Advancing Academic Research through Innovation in eScience and Data Science Technologies, the Netherlands e-Science Center and The Dutch Research Council: 2020
 - ETEC Proposal Review, the Netherlands e-Science Center (NLeSC) and The Dutch Research Council (NWO): 2020

Awarded Proposals

9/2020 ~ 8/2025	RAPIDS2: A SciDAC Institute for Computer Science, Data, and Artificial Intelligence, DoE, SciDAC Institute for Computer Science and Data, Senior Personnel
12/2019 ~ 11/2022	Brisbane: Productive Programming Systems in the Era of Extremely Heterogeneous and Ephemeral Computer Architectures, DoD, Advanced Computing Initiative (ACI), \$2.25M, Senior Personnel
7/2018 ~ 6/2022	COSMIC CASTLE, DoD, Defense Advanced Research Projects Agency (DARPA), \$6M, Senior Personnel
9/2017 ~ 8/2020	RAPIDS: a SciDAC Institute for Resource and Application Productivity through computation, Information, and Data Science, DoE, SciDAC Institute for Computer Science and Data, \$18M, Senior Personnel
10/2016 ~ 9/2019	PROTEAS: PROgramming Toolchain for Emerging Architectures and System, DoE, Exascale Computing Project (ECP), \$3.8M, Senior Personnel
10/2016 ~ 9/2018	Tahoe: Designing and Programming Exascale Memory Hierarchies, DoE, Laboratory Directed Research & Development Program, LDRD Project, \$1M, Co-Principle Investigator
4/2016 ~ 3/2017	Understanding the Interface Driven Magnetic Properties of Topological Insulators Using a GPU Accelerated First-Principles All-Electron Code, DOE, Laboratory Directed Research & Development Program, SEED Project, \$95K, Co-Principal Investigator
10/2015 ~ 9/2016	Electronic Structure Based Discovery of Hybrid Photovoltaic Materials on Next-Generation HPC Platforms, DOE, ALCF Theta Early Science Program, Tier 2 Early Science Project, Co-Principal Investigator

- 7/2014 ~ 6/2017 ARES: Abstract Representations for the Extreme-Scale Stack, DOE, Office of Advanced Scientific Computing Research, \$3M, Senior Personnel
- 10/2013 ~ 9/2016 Vancouver2: Improving Programmability of Contemporary Heterogeneous Architectures, DOE, Office of Advanced Scientific Computing Research, \$2.1M, Senior Personnel
- 7/2013 ~ 12/2014 Programmer-Guided Reliability and Trade-Offs with Energy and Performance, DOD, Department of Defense Advanced Computing Initiative (ACI) Fiscal Year 2012, \$1.11M, Senior Personnel

Publications

Anthony M. Cabrera, Aaron R. Young, Jacob Lambert, Zhili Xiao, Amy An, **Seyong Lee**, Zheming Jin, Jungwon Kim, Jeremy Buhler, Roger D. Chamberlain, and Jeffrey S. Vetter, Towards Evaluating High-Level Synthesis Portability and Performance Between Intel and Xilinx FPGAs, *9th International Workshop on OpenCL and SYCL (IW'OCCL)*, 2021

Gregory Herschlag, **Seyong Lee**, Jeffrey S. Vetter, and Amanda Randles, Analysis of GPU Data Access Patterns on Complex Geometries for the D3Q19 Lattice Boltzmann Algorithm, *Transactions on Parallel and Distributed Systems (TPDS)*, 2021

Camille Coti, Joel E. Denny, Kevin Huck, **Seyong Lee**, Allen D. Malony, Sameer Shende, and Jeffrey S. Vetter, OpenACC Profiling Support for Clang and LLVM using Clacc and TAU, *Workshop on Programming and Performance Visualization Tools (ProTools 20)*, in conjunction with SC20, 2020

Mohammad Alaul Haque Monil, **Seyong Lee**, Jeffrey S. Vetter, and Allen D. Malony, Understanding the Impact of Memory Access Patterns in Intel Processors, *MCHPC'20: Workshop on Memory Centric High Performance Computing*, in conjunction with SC20, 2020.

Mohammad Alaul Haque Monil, Mehmet E. Belviranli, **Seyong Lee**, Jeffrey S. Vetter, and Allen D. Malony, MEPHESTO: Modeling Energy-Performance in Heterogeneous SOCs and Their Trade-Offs, *The International Conference on Parallel Architectures and Compilation Techniques (PACT)*, 2020.

Jacob Lambert, **Seyong Lee**, Jeffrey S. Vetter, and Allen D. Malony, CCAMP: An Integrated Translation and Optimization Framework for OpenACC and OpenMP, *SC 2020: The International Conference for High Performance Computing, Networking, Storage, and Analysis*, 2020.

Ryuta Tsunashima, Ryohei Kobayashi, Norihisa Fujita, Taisuke Boku, **Seyong Lee**, Jeffrey Vetter, Hitoshi Murai, Masahiro Nakao and Mitsuhsa Sato, OpenACC unified programming environment for GPU and FPGA multi-hybrid acceleration, *13th International Symposium on High-level Parallel Programming and Applications (HLPP)*, Porto, Portugal, July 9-10, 2020.

Jacob Lambert, **Seyong Lee**, Jeffrey S. Vetter, and Allen D. Malony, In-Depth Optimization with the OpenACC-to-FPGA Framework on an Arria 10 FPGA, *The Ninth International Workshop on Accelerators and Hybrid Exascale Systems (AsHES)*, in conjunction with IPDPS20, New Orleans, LA, USA, 2020.

Roberto Gioiosa, Burcu O. Mutlu, **Seyong Lee**, Jeffrey S. Vetter, Giulio Picierro, and Marco Cesati, The Minos Computing Library: Efficient Parallel Programming for Extremely Heterogeneous Systems, *Proceedings of the 13th Annual Workshop on General Purpose Processing using Graphics Processing Unit (GPGPU'20)*, 2020.

Blaise Tine, Fares Elsabbagh, **Seyong Lee**, Jeffrey Vetter, and Hyesoon Kim, Cash: A Single-Source Hardware-Software Codesign Framework for Rapid Prototyping, *28th ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA 2020)*, Poster, Seaside, California, USA, 2020.

Blaise Tine, **Seyong Lee**, Jeffrey Vetter, and Hyesoon Kim, Productive Hardware Designs using

Hybrid HLS-RTL Development, *28th ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA 2020)*, Poster, Seaside, California, USA, 2020.

Ryuta Tsunashima, Ryohei Kobayashi, Norihisa Fujita, Ayumi Nakamichi, Taisuke Boku, **Seyong Lee**, Jeffrey Vetter, Hitoshi Murai, and Mitsuhsa Sato, Enabling OpenACC Programming on Multi-hybrid Accelerated with GPU and FPGA, *International Conference on High Performance Computing in Asia-Pacific Region (HPC Asia 2020)*, Poster, Fukuoka, Japan, 2020.

Forrest Shriver, **Seyong Lee**, Steven Hamilton, Justin Watson and Jeffrey Vetter, Enhancing Monte Carlo proxy applications on GPUs, *10th IEEE International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS19), in conjunction with SC19*, 2019.

Forrest Shriver, **Seyong Lee**, Steven Hamilton, Jeffrey Vetter, and Justin Watson, VEXS, An Open Platform for the Study of Continuous-Energy Neutron Transport Cross-Section Lookup Algorithms on GPUs, *MC19: International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering*, 2019.

David Ojika, Ann Gordon-Ross, Herman Lam, Shinjae Yoo, Younggang Cui, Zhihua Dong, Kirstin Kleese Van Dam, **Seyong Lee**, and Thorsten Kurth, PCS: A Productive Computational Science Platform, *International Workshop on Exploitation of High Performance Heterogeneous Architectures and Accelerators (WEHA), in conjunction with HPCS19*, 2019.

Jacob Lambert, **Seyong Lee**, Allen D. Malony, and Jeffrey S. Vetter, CCAMP: OpenMP and OpenACC Interoperable Framework, *Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms (HeteroPar), in conjunction with Euro-Par19*, 2019.

Seyong Lee, John Gounley, Amanda Randles, and Jeffrey S. Vetter, Performance Portability Study for Massively Parallel Computational Fluid Dynamics Application on Scalable Heterogeneous Architectures, *Journal of Parallel and Distributed Computing (JPDC) volume 129, Pages 1-13*, July 2019.

Joel E. Denny, **Seyong Lee**, and Jeffrey S. Vetter, Clacc: Translating OpenACC to OpenMP in Clang, *IEEE/ACM 5th Workshop on the LLVM Compiler Infrastructure in HPC (LLVM-HPC), in conjunction with SC18*, 2018.

Mehmet E. Belviranlı, **Seyong Lee**, and Jeffrey S. Vetter, Programming the EMU Architecture: Algorithm Design Considerations for Migratory-threads-based Systems, *SC 2018: The International Conference for High Performance Computing, Networking, Storage, and Analysis*, Poster, 2018.

Seyong Lee, Jacob Lambert, Jungwon Kim, Jeffrey S. Vetter, and Allen D. Malony, OpenACC to FPGA: A Directive-Based High-Level Programming Framework for High-Performance Reconfigurable Computing, *SC 2018: The International Conference for High Performance Computing, Networking, Storage, and Analysis*, Poster, 2018.

Pak Markthub, Mehmet E. Belviranlı, **Seyong Lee**, Jeffrey S. Vetter, and Satoshi Matsuoka, DRAGON: Breaking GPU Memory Capacity Limits with Direct NVM Access, *SC 2018: The International Conference for High Performance Computing, Networking, Storage, and Analysis*, November 2018.

Michael Wolfe, **Seyong Lee**, Jungwon Kim, Xiaonan Tian, Rengan Xu, Barbara Chapman, Sunita Chandrasekaran, The OpenACC data model: Preliminary study on its major challenges and implementations, *Parallel Computing: systems & applications Volume 78, Pages 15-27*, October 2018.

Mehmet E. Belviranlı, **Seyong Lee**, and Jeffrey S. Vetter, Designing Algorithms for the EMU Migrating-threads-based Algorithms, *HPEC18: IEEE High Performance Extreme Computing Conference, Best Paper Finalist*, September 2018.

Anshu Dubey, Samuel Williams, Phillip C. Roth, **Seyong Lee**, Tom Peterka, Hongzhang Shan,

- Chad Steed, RAPIDS: Applications Engagement and Community Outreach, *SciDAC4 RAPIDS PI Meeting*, Poster, July 2018.
- Jeffrey S. Vetter, Paul Hovland, Samuel Williams, Costin Iancu, **Seyong Lee**, Swann Pereneau, Phillip C. Roth, Kevin Huck, RAPIDS: Platform Readiness, *SciDAC4 RAPIDS PI Meeting*, Poster, July 2018.
- Ivy Bo Peng, Jeffrey S. Vetter, Shirley V. Moore, and **Seyong Lee**, Exploring Tuyere: Enabling Scalable Memory Workloads for System Exploration, *HPDC18: The 27th International Symposium on High-Performance Parallel and Distributed Computing*, June 2018.
- Jacob B. Lambert, **Seyong Lee**, Jungwon Kim, Jeffrey S. Vetter, and Allen D. Malony, Directive-based, High-Level Programming and Optimizations for High-Performance Computing with FPGAs, *ICS 2018: The 32nd ACM International Conference on Supercomputing*, June 2018.
- Pak Markthub, Mehmet E. Belviranlı, **Seyong Lee**, Jeffrey S. Vetter, and Satoshi Matsuoka, Efficiently Enlarging GPU Memory Capacity with NVM, *GTC'18: GPU Technology Conference*, Poster, March 2018.
- Mehmet E. Belviranlı, **Seyong Lee**, Jeffrey S. Vetter, and Laxmi N. Bhuyan, Juggler: A Dependency-Aware Task Based Execution Framework for GPUs, *PPoPP18: ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, 2018.
- Gregory Herschlag, Amanda Randles, **Seyong Lee**, and Jeffrey S. Vetter, GPU Data Access on Complex Geometries for D3Q19 Lattice Boltzmann Method, *IPDPS18: IEEE International Parallel and Distributed Processing Symposium*, 2018.
- Kaixi Hou, Hao Wang, Wu-chun Feng, Jeffrey S. Vetter, and **Seyong Lee**, Highly Efficient Compensation-based Parallelism for Wavefront Loops on GPUs, *IPDPS18: IEEE International Parallel and Distributed Processing Symposium*, 2018.
- Seyong Lee**, A Framework for Directive-Based High-Performance Reconfigurable Computing, *OpenACC for Programmers: Concepts and Strategies*, Book Chapter, 2017.
- Jungwon Kim, **Seyong Lee**, and Jeffrey S. Vetter, PapyrusKV: A High-Performance Parallel Key-Value Store for Distributed NVM Architectures, *SC'17: ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis*, November 2017.
- Michael Wolfe, **Seyong Lee**, Jungwon Kim, Xiaonan Tian, Rengan Xu, Sunita Chandrasekaran, and Barbara Chapman, Implementing the OpenACC Data Model, *The Seventh International Workshop on Accelerators and Hybrid Exascale Systems (AsHES) in conjunction with IPDPS17*, 2017.
- Joel E. Denny, **Seyong Lee**, and Jeffrey S. Vetter, Language-Based Optimizations for Persistence on Nonvolatile Main Memory Systems, *31th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, 2017.
- Jungwon Kim, Kittisak Sajjapongse, **Seyong Lee**, and Jeffrey S. Vetter, Design and Implementation of Papyrus: Parallel Aggregate Persistent Storage, *31th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, 2017.
- Joel E. Denny, **Seyong Lee**, and Jeffrey S. Vetter, NVL-C: Static Analysis Techniques for Efficient, Correct Programming of Non-Volatile Main Memory Systems, *HPDC'16: Proceedings of the ACM Symposium on High-Performance and Distributed Computing*, 2016.
- Jungwon Kim, **Seyong Lee**, and Jeffrey S. Vetter, IMPACC: A Tightly Integrated MPI+OpenACC Framework Exploiting Shared Memory Parallelism, *HPDC'16: Proceedings of the ACM Symposium on High-Performance and Distributed Computing*, 2016.
- Seyong Lee**, Jungwon Kim, and Jeffrey S. Vetter, OpenACC to FPGA: A Framework for

- Directive-Based High-Performance Reconfigurable Computing, *30th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, 2016.
- David E. Bernholdt, Wael R. Elwasif, Christos Kartaklis, **Seyong Lee**, and Tiffany M. Mintz, Programmer-Guided Reliability for Extreme-Scale Applications, *International Journal of High Performance Computing Applications*, 2016.
- Jacob Lambert, **Seyong Lee**, Jungwon Kim, and Jeffrey S. Vetter, OpenACC to FPGA: Optimization of Directive-Based Programming for Reconfigurable Devices, *ORISE Summer 2016 Graduate, Post Graduate, Employee Participant, and Faculty Poster Session*, Poster, 2016.
- Joel E. Denny, **Seyong Lee**, and Jeffrey S. Vetter, FITL: Extending LLVM for the transaction of fault-injection directives, *the Second Workshop on the LLVM Compiler Infrastructure in HPC (LLVM'15) in conjunction with SC*, ACM, 2015.
- David E. Bernholdt, Wael R. Elwasif, Christos Kartaklis, **Seyong Lee**, and Tiffany M. Mintz, Programmer-Guided Reliability for Extreme-Scale Applications, *1st International Workshop on Fault Tolerant Systems (FTS) in conjunction with IEEE Cluster*, 2015.
- Amit Sabne, Putt Sakdhnagool, **Seyong Lee**, and Jeffrey S. Vetter, Understanding Portability of a High-level Programming Model on Contemporary Heterogeneous Architectures, *IEEE Micro*, 2015.
- Seyong Lee**, Jeremy S. Meredith, and Jeffrey S. Vetter, COMPASS: A Framework for Automated Performance Modeling and Prediction, *ACM International Conference on Supercomputing (ICS)* 2015.
- Jungwon Kim, **Seyong Lee**, and Jeffrey S. Vetter, An OpenACC-based Unified Programming Model for Multi-accelerator Systems, *Proceedings of the 20th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP)*, Poster, 2015.
- Sarat Sreepathi, Megan Grodowitz, Robert Lim, Philip Taffet, Philip Roth, Jeremy Meredith, **Seyong Lee**, Dong Li, and Jeffrey S. Vetter, Application Characterization using Oxbow Toolkit and PADS Infrastructure, *First International Workshop on Hardware-Software Co-Design for High Performance Computing (Co-HPC) in conjunction with SC14*, 2014.
- Seyong Lee** and Jeffrey S. Vetter, OpenARC: Extensible OpenACC Compiler Framework for Directive-Based Accelerator Programming Study, *Workshop on Accelerator Programming Using Directives (WACCPD) in conjunction with SC14*, 2014.
- Amit Sabne, Putt Sakdhnagool, **Seyong Lee**, and Jeffrey S. Vetter, Evaluating Performance Portability of OpenACC, *LCPC'14: The 27th International Workshop on Languages and Compilers for Parallel Computing*, 2014.
- Seyong Lee** and Jeffrey S. Vetter, OpenARC: Open Accelerator Research Compiler for Directive-Based, Efficient Heterogeneous Computing, *HPDC'14: Proceedings of the ACM Symposium on High-Performance and Distributed Computing*, Short Paper, 2014.
- Seyong Lee**, Dong Li, and Jeffrey S. Vetter, Interactive Program Debugging and Optimization for Directive-Based, Efficient GPU Computing, *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2014.
- Seyong Lee** and Jeffrey S. Vetter, OpenARC: Open Accelerator Research Compiler for Directive-Based, Heterogeneous Computing, *GTC'14: GPU Technology Conference*, Poster, 2014.
- Dong Li, **Seyong Lee**, and Jeffrey S. Vetter, Evaluate the Viability of Application-Driven Cooperative CPU/GPU Fault Detection, *Workshop on Resiliency in High Performance Computing (Resilienc)*, in conjunction with Euro-Par, 2013.
- Jeffrey S. Vetter, **Seyong Lee**, Dong Li, Gabriel Marin, Jeremy Meredith, Philip C. Roth, and Kyle

Spafford, Quantifying Architectural Requirements of Contemporary Extreme-Scale Scientific Applications, *International Workshop on performance modeling, benchmarking and simulation of high performance computer systems (PMBS) in conjunction with SC13*, 2013.

Seyong Lee and Rudolf Eigenmann, OpenMPC: Extended OpenMP for Efficient Programming and Tuning on GPUs, *International Journal of Computational Science and Engineering Volume 8, Issue 1*, February 2013.

Faraz Ahmad, **Seyong Lee**, Mithuna Thottethodi, and T. N. VijayKumar, MapReduce with Communication Overlap (MaRCO), *Journal of Parallel and Distributed Computing (JPDC), Volume 73, Issue 5*, May 2013.

Seyong Lee and Jeffrey S. Vetter, Early Evaluation of Directive-Based GPU Programming Models for Productive Exascale Computing, *SC'12: ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis*, November 2012.

Seyong Lee, Toward Compiler-Driven Adaptive Execution, *LAMBERT Academic Publishing*, ISBN 978-3-659-20487-6, 2012.

Kyle Spafford, Jeremy S. Meredith, **Seyong Lee**, Dong Li, Philip C. Roth, and Jeffrey S. Vetter, The Tradeoffs of Fused Memory Hierarchies in Heterogeneous Computing Architectures, *the Proceedings of the ACM International Conference on Computing Frontiers*, May 2012.

Seyong Lee and Jeffrey S. Vetter, Moving Heterogeneous GPU Computing into the Mainstream with Directive-Based, High-Level Programming Models (position paper), *DOE Exascale Research Conference*, April 2012.

Faraz Ahmad, **Seyong Lee**, Mithuna Thottethodi, and T.N. VijayKumar, PUMA: Purdue MapReduce Benchmarks Suite, *ECE Technical Reports, Purdue University*, October 2012.

Seyong Lee, Toward Compiler-Driven Adaptive Execution and Its Application to GPU Architectures, PhD Thesis, *School of Electrical and Computer Engineering – Purdue University*, May 2011.

Seyong Lee and Rudolf Eigenmann, OpenMPC: Extended OpenMP Programming and Tuning for GPUs, *SC'10: ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (Best Student Paper Award)*, November 2010.

Chirag Dave, Hansang Bae, Seung-Jai Min, **Seyong Lee**, Rudolf Eigenmann, and Samuel Midkiff, Cetus: A source-to-Source Compiler Infrastructure for Multicores, *IEEE Computer Volume 42, Issue 12, pp36-42*, December 2009.

Seyong Lee, Seung-Jai Min, and Rudolf Eigenmann, OpenMP to GPGPU: A Compiler Framework for Automatic Translation and Optimization, *Symposium on Principles and Practice of Parallel Programming (PPoPP)*, February 2009.

Hansang Bae, Leonardo Bacheга, Chirag Dave, Sang-Ik Lee, **Seyong Lee**, Seung-Jai Min, Rudolf Eigenmann, and Samuel Midkiff, Cetus: A Source-to-Source Compile Infrastructure for Multicore, *14th Workshop on Compilers for Parallel Computing (CPC)*, January 2009.

Seyong Lee and Rudolf Eigenmann, Adaptive Runtime Tuning of Parallel Sparse Matrix-Vector Multiplication on Distributed Memory Systems, *22nd ACM International Conference on Supercomputing (ICS)*, June 2008.

Seyong Lee and Rudolf Eigenmann, Adaptive Tuning in a Dynamically Changing Resource Environment, *Workshop on National Science Foundation Next Generation Software Program (NSFNGS) held in conjunction with the IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, April 2008.

Seyong Lee, Xiaojuan Ren, and Rudolf Eigenmann, Efficient Content Search in iShare, a P2P

based Internet-Sharing System, *2nd Workshop on Large-scale, volatile Desktop Grids (PCGrid) held in conjunction with the IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, April 2008.

Faraz Ahmad, **Seyong Lee**, Mithuna Thottethodi, and T. N. VijayKumar, MapReduce with Communication Overlap (MaRCO), ECE Technical Reports TR-ECE-11-07, Electrical and Computer Engineering, Purdue University, November 2007.

Xiaojuan Ren, **Seyong Lee**, Rudolf Eigenmann, and Saurabh Bagchi, Prediction of Resource Availability in Fine-Grained Cycle Sharing Systems and Empirical Evaluation, *Journal of Grid Computing Volume 5, Number 2, pp173-195*, June 2007.

Xiaojuan Ren, **Seyong Lee**, Rudolf Eigenmann, and Saurabh Bagchi, Resource Failure Prediction in Fine-Grained Cycle Sharing Systems, *The 15th IEEE International Symposium on High Performance Distributed Computing (Nominated for Best Paper Award)*, June 2006.

Xiaojuan Ren, **Seyong Lee**, Saurabh Bagchi, and Rudolf Eigenmann, Resource Fault Prediction in Fine-Grained Cycle Sharing, *DSN-2005: The International Conference on Dependable Systems and Networks*, Fast Abstracts, June 2005.

Invited and Contributed Presentations

Seyong Lee, Performance Portability in the Era of Extremely Heterogeneous Computing, KIISE Computer System Society, South Korea, February 2021.

Seyong Lee, Introduction to Exascale Computing Project, *Samsung Advanced Institute of Technology (SAIT)*, South Korea, January 2021.

Seyong Lee, ETRI/ORNL Supercomputing Meeting, *Electronics and Telecommunications Research Institute (ETRI)*, South Korea, January 2021.

Seyong Lee, Challenges and Opportunities for AI-Driven, Directive-based Heterogeneous Computing, invited talk, *Workshop on Program Synthesis for Scientific Computing (PSW)*, August 2020.

Seyong Lee and Jungwon Kim, Directive-based High-Level Porting of GR Blocks onto Heterogeneous Architectures, *Gnuradio 4.0 Workshop*, July 2020.

Jeffrey S. Vetter and **Seyong Lee**, Striving for Performance Portability of Software Radio Software in the Era of Heterogeneous SoCs, invited talk, *Free Open Source Developers' European Meeting (FOSDEM)*, Brussels, Belgium, February 2020.

Seyong Lee, Parallel Programming with OpenACC, tutorial, *Korean Supercomputing Conference (KSC)*, Seoul, South Korea, September 2019.

Seyong Lee, Advanced OpenACC Programming, tutorial, *Korea Institute of Science and Technology Information (KISTI)*, Daejeon, South Korea, September 2019.

Seyong Lee, Exascale Computing Project, invited talk, *Korea Institute of Science and Technology information (KISTI)*, Daejeon, South Korea, September 2019.

Oscar Hernandez, **Seyong Lee**, and Sunita Chandrasekaran, On-going Efforts with Open Source Compilers, invited talk, *Annual OpenACC Meeting*, Kobe, Japan, September 2019.

Seyong Lee, OpenARC: Extensible Compiler Framework for Directive-Based, Efficient Heterogeneous Computing, invited talk, *Electrical and Computer Engineering Seminar, University of Delaware*, March 2019.

Seyong Lee, Array Reduction, *OpenACC Technical Meeting*, May 2018.

Seyong Lee, OpenACC Status Update: Recent Highlights and New Features in OpenACC 2.6, invited talk, *SciDAC4 RAPIDS Platform Readiness Monthly Meeting*, March 2018.

Jeffrey S. Vetter and **Seyong Lee**, Working Toward Performance Portability for FPGAs in High Performance Computing, invited talk, *3rd International Workshop on FPGA for HPC (IWFH)*, Tokyo, Japan, March 2018.

Shirley V. Moore, Jeffrey S. Vetter, **Seyong Lee**, and Ivy Peng, Exploring Extreme Heterogeneity with Analytical Modeling, invited talk, *SLAM Conference on Parallel Processing for Scientific Computing (SLAMPP)*, Tokyo, Japan, March 2018.

Seyong Lee, Jungwon Kim, and Jeffrey S. Vetter, OpenARC: A Compiler Framework for Directive-Based High Performance Reconfigurable Computing, invited talk, *SLAM Conference on Parallel Processing for Scientific Computing (SLAMPP)*, Tokyo, Japan, March 2018.

Seyong Lee, Jungwon Kim, and Jeffrey S. Vetter, OpenACC to FPGA: Directive-Based High Performance Reconfigurable Computing, invited talk, *BNL FPGA ML Working Group Meeting*, Upton, NY, January 2018.

Seyong Lee, Joel E. Denny, Jungwon Kim, Mehmet E. Belviranli and Jeffrey S. Vetter, OpenARC: Extensible Compiler Framework for Directive-Based, Efficient Heterogeneous Computing Study, invited talk, *JST/CREST International Symposium on Post Petascale System Software (ISP2S2)*, Tokyo, Japan, December 2017.

Seyong Lee, Jungwon Kim, and Jeffrey S. Vetter, OpenACC to FPGA: A Framework for Directive-Based High-Performance Reconfigurable Computing, paper presentation, *30th IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, May 2016.

Seyong Lee, OpenACC and Memory Hierarchies, *OpenACC Technical Meeting*, June 2016.

Seyong Lee, Jeremy S. Meredith, and Jeffrey S. Vetter, COMPASS: A Framework for Automated Performance Modeling and Prediction, paper presentation, *ACM International Conference on Supercomputing (ICS)*, June 2015.

Seyong Lee and Jeffrey S. Vetter, Extended OpenACC Programming to Exploit GPU-Specific Features Still at a High Level, invited talk, *GPU Technology Conference (GTC)*, March 2015.

Seyong Lee and Jeffrey S. Vetter, OpenARC: Extensible OpenACC Compiler Framework for Directive-Based Accelerator Programming Study, paper presentation, *Workshop on Accelerator Programming Using Directives (WACCPD) in conjunction to SC14*, November 2014.

Seyong Lee, OpenARC: Open Accelerator Research Compiler, invited talk, *OpenARC workshop*, May and August 2014.

Seyong Lee, Dong Li, and Jeffrey S. Vetter, Interactive Program Debugging and Optimization for Directive-Based, Efficient GPU Computing, paper presentation, *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, May 2014.

Seyong Lee, OpenARC: Open Accelerator Research Compiler, invited talk, *OpenACC Technical Forums Face-to-Face Meeting*, June 2013.

Seyong Lee, Early Evaluation of Directive-Based GPU Programming Models, invited talk, *the first KIISE-KOCSEA HPC SIG Joint Workshop on High Performance and Throughput Computing*, November 2012

Seyong Lee and Jeffrey S. Vetter, Early Evaluation of Directive-Based GPU Programming Models for Productive Exascale Computing, paper presentation, *SC'12: ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis*, November 2012.

Seyong Lee, Directive-Based GPU Programming Models: Current Status and Tuning Opportunities, invited talk, *CScADS Summer Workshop on Libraries and Autotuning for Extreme-scale*

Applications, August 2012.

Seyong Lee, Toward Compiler-Driven Adaptive Execution, invited talk, *Future Technologies Group Colloquium*, Oak Ridge National Laboratory, January 2011.

Seyong Lee and Rudolf Eigenmann, OpenMPC: Extended OpenMP Programming and Tuning for GPUs, paper presentation, *SC'10: ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis*, November 2010.

Seyong Lee, Seung-Jai Min, and Rudolf Eigenmann, OpenMP to GPGPU: A Compiler Framework for Automatic Translation and Optimization, paper presentation, *Symposium on Principles and Practice of Parallel Programming (PPoPP)*, February 2009.

Seyong Lee, Xiaojuan Ren, and Rudolf Eigenmann, Efficient Content Search in iShare, a P2P based Internet-Sharing System, paper presentation, *2nd Workshop on Large-scale, volatile Desktop Grids (PCGrid) held in conjunction with the IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, April 2008.

Patents

“Algorithm and Hardware Architecture of Multi-channel FSK Modem for Powerline Communication”, Author: Jintae Kim, Jihyoun Kim, Taesang Yoo, and **Seyong Lee**

Honors and Awards

IEEE-CS TCHPC Award for Excellence for Early Career Researchers in High Performance Computing

Awarded up to 3 individuals who have made outstanding, influential, and potentially long-lasting contributions in the field of high performance computing within 5 years of receiving their PhD degrees (<https://www.computer.org/web/pressroom/TCHPC-Award-2016>).

Best Student Paper Award in Proceedings of the 2010 ACM/IEEE conference on Supercomputing (SC10)

The paper on the high-level GPU programming (“OpenMP to GPGPU: A Compiler Framework for Automatic Translation and Optimization”) was selected as the most cited paper among all papers published in the Symposium on Principles and Practices of Parallel Programming (PPOPP) between 2009 and 2014 (cited 541 times as of February 2019).

The Samsung Lee Kun Hee Scholarship Foundation (2004~2008)

Awarded to 50 B.S. students, 25 M.S. students, and 25 PhD. Students in all area with focus on science and engineering area
Full-tuition and living expense for four years

IT Scholarship of Ministry of Information and Communication Republic of Korea (2002, 2003)

Awarded to 20 M.S. students and 50 PhD. students in IT' area through highly competitive selection procedure

Korea Foundation for Advanced Studies (KFAS) College Student Scholarship (1997, 1998)

Awarded to 20 students in EECS through highly competitive selection procedures

Chungbuk Association College Student Scholarship(1995~1998)

Awarded to top 5 all high school graduates in Chungbuk Province
Full-tuition for four years

Ranked 50th of all applicants at the Korea National College Entrance Exam (1994)

Ranked 50th of all applicants in South Korea (50/757,488)

Skills

Programming & Tools Experience

- GPU programming using high-level directives (PGI Accelerator, HMPP, R-Stream, OpenACC, OpenMPC)
- Parallel programming and performance tuning using MPI, OpenMP, CUDA, and OpenCL
- Simulation and Analysis using C++ and MATLAB
- Developed a compiler system for source-to-source transformation and optimizations, which is written in Java.
- Optimized the performance of the Apache Hadoop MapReduce System and Distributed File System (DFS), which are written in Java.
- Programming Experience with C/C++, Java, Fortran, Python, Perl, Tcl, PHP, SQL, and shell-script languages
- Programming experience on various Unix/Linux environments such as RHEL, Debian, Ubuntu, and Solaris.
- Programming experience on RDBMS such as Oracle and MySQL
- Experience on Internet-sharing/Cloud computing middlewares, such as HUBzero and iShare.
- Experience on Content Management Systems (Joomla and Expression Engine) and Rappture Toolkit,
- Experience on Apache HTTP server, Java Application Server (JBoss), and Java EE applications
- Emulation with FPGA using ALTERA MAX+II and Xilinx Foundation
- ASIC Design and Simulation using HDL (VHDL, Verilog) and tool (Synopsis, VerilogXL)