

Seyong Lee

121 Derby Run Drive,
Farragut, TN 37934
Phone : 765-532-0877
Email : lees2@ornl.gov

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

Major Courses Taken

ECE 573 – Compilers and Translator Writing System
ECE 663 – Compiler Code Generation, Optimization, and Parallelization
ECE 563 – Programming Parallel Machines
ECE 565 – Computer Architecture
ECE 666 – Advanced Computer Systems (Parallel Computer Architecture)
ECE 572 – Fault-Tolerant Computer Systems
ECE 673 – Distributed Computing System
ECE 608 – Computational Models and Methods
ECE 547 – Introduction to Computer Communication Networks
ECE 600 – Random Variables
CS 503 – Operating Systems
CS 590I – Information Retrieval
MA 575 – Graph Theory
MA 518 – Advanced Discrete Mathematics

Research Experience

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

5/2011 ~ present

Computer Scientist, Future Technology Group, Oak Ridge National Laboratory

- Develop high-level programming models for future, heterogeneous computing systems.

9/2009 ~ 12/2009

Software Engineer (Intern), NEEScomm, Discovery Park, Purdue University

- 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. (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 managing jobs such as account managing, CVS and other utility environment setup, and etc.

Professional Service

Reviewing (Conferences, Workshops, Journals, and research proposals)

- 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
- International Conference on Supercomputing (ICS):
 - 2008 and 2011
- ACM/IEEE International Conference for High Performance Computing, Networking, Storage, and Analysis (SC):
 - 2007
- 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
- Journal of Parallel and Distributed Computing (JPDC):
 - 2009
- International Journal of High Performance Computing:
 - 2012
- International Journal of Software and Systems Modeling (SOSYM):
 - 2011
- Journal of Software: Practice and Experience (SPE):
 - 2010
- Workshop on Languages and Compilers for a Parallel Computing (LCPC):
 - 2006, 2007, and 2011
- 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
- The General Research Fund, the Research Grants Council of Hong Kong:
 - 2011

Member of the OpenMP Language Committee and OpenMP Accelerator Subcommittee

Guest Editor for the special issue on “Advanced Development in Software of GPU Cluster Computing” in Journal of Advances in Software Engineering

Publications

Seyng 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 and Rudolf Eigenmann, OpenMPC: Extended OpenMP for Efficient Programming and Tuning on GPUs, *International Journal of Computational Science and Engineering*, 2012

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

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 Bacheaga, 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, 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, talk, *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, talk, *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, talk, *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, Jiyoung Kim, Taesang Yoo, and **Seyong Lee**

Honors and Awards

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

Awarded to 50 B.S. students, 25 M.S. students, and 25 Ph.D. 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 Ph.D. 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, and CUDA
- 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)