

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

- 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.
 - 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

5/2011 ~ present

Computer Scientist, Future Technology Group, Oak Ridge National Laboratory (<http://ft.ornl.gov>)

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

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

Program Committee Member

- The IEEE International Conference on Parallel and Distributed Systems (ICPADS): 2013, 2014, 2015, and 2016
- ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP): 2014 (external review committee member)
- IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing (CCGrid): 2015 and 2016 (Technical Poster Committee)
- International Conference on Parallel Processing (ICPP): 2013
- International Symposium on Computing and Networking (CANDAR): 2016
- International Workshop on Programming Models, Languages, and Compilers for Manycore and Heterogeneous Architectures (PLC): 2015
- International Workshop on Representative Applications (WRAP): 2015
- International Workshop on Accelerator Programming Using Directives (WACCPD): 2014, 2015, 2016
- International Workshop on Accelerators and Hybrid Exascale Systems (AsHES): 2016
- International Workshop on Legacy HPC Application Migration (LHAM): 2016

Member of OpenMP Language Committee and OpenMP Accelerator Subcommittee

Member of the OpenACC Technical Committee and Test-Suite Committee

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

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

External Reviewer (Journals, Conferences, Workshops, and research proposals)

Journals

- International Journal of High Performance Computing (IJHPC): 2015 and 2016
- IEEE Transactions on Parallel and Distributed Systems (TPDS): 2014 and 2016
- ACM Transactions on Modeling and Performance Evaluation of Computing Systems (ToMPECS): 2015
- International Journal on Parallel Computing (ParCo): 2013 and 2015
- International Journal of Computing Science and Applications (CyS): 2015
- Journal of Parallel and Distributed Computing (JPDC): 2009
- International Journal of High Performance Computing (IJHPC): 2012
- ACM Transactions on Architecture and Code Optimization (ACMTACO): 2013 and 2014
- International Journal of Software and Systems Modeling (SOSYM): 2011
- Journal of Software: Practice and Experience (SPE): 2010

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

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

Awarded Proposals

- 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

- 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
- 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 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 Bachega, 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, 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)*, 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)* 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)*, 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*, 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)*, 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, *CS&ADS 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 300 times as of April, 2014).

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 Scholarip(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)