Hyojin Sung

hsung@postech.ac.kr ◇ hyojin.sung@gmail.com

EDUCATION

University of Illinois at Urbana-Champaign *August 2015*
Ph.D. in Computer Science
Overall GPA: 4.0/4.0
Doctoral Thesis: ”DeNovo: Rethinking the Memory Hierarchy for Disciplined Parallelism” , advised by Prof. Sarita V. Adve

University of California, San Diego *August 2008*
M.S. in Computer Science and Engineering
M.S. Thesis: ”A Portable MATLAB Front-End for Tiled Microprocessors”, advised by Prof. Michael B. Taylor

Seoul National University *June 2004*
B.S. in Computer Science and Engineering
B.S. Thesis: ”Survey on Code-Size Optimization Techniques in Embedded Systems”, advised by Prof. Sang-Lyul Min

Seoul National University *February 2000*
B.A. in English Literature

HONORS AND AWARDS

ACM Doctoral Dissertation Award Nomination, 2015
– One of the two nominations from the Department of Computer Science at UIUC.

Runner-up for David J. Kuck Outstanding Ph.D Thesis Award, 2015
– Department of Computer Science at UIUC.

Andrew and Shana Laursen Fellowship, 2014-15
– Department of Computer Science at UIUC.

W. J. Poppelbaum Memorial Award, 2014
– Department of Computer Science at UIUC. This award is intended for graduate students in computer hardware or architecture, based on academic merit and creativity.

IEEE Micro Top Picks from the Computer Architecture Conferences, May/June 2014
– Recognizes the most significant research papers published in computer architecture in 2013 based on novelty and long-term impact. Only 12 papers were selected.

The 2nd Rising Stars in EECS workshop at MIT, 2013
– Invited top female PhDs and postdocs for scientific discussions and informal sessions on an academic career.

Best paper award in PACT, 2011
– 20th International Conference on Parallel Architectures and Compilation Techniques (PACT)

The Feng Chen memorial award, 2011 and 2014
– Department of Computer Science at UIUC, for the best paper award in PACT 2011 and the IEEE Micro Top Picks in 2014

Merit scholarships, Spring 2002 to Fall 2003
– Department of Computer Science and Engineering at Seoul National University

Merit scholarships, Spring 1997 to Fall 1999
– Department of English Literature at Seoul National University

EXPERIENCE

POSTECH (Pohang University of Science and Technology) July 2019 – Present
Asssitant Professor *Pohang, South Korea*

- Auto-tuning for deep learning compilers
- Compiler support and optimizations for Processor-in-Memory hardware
- OpenCL compiler and runtime for AB21 supercomputers

IBM Corporation May 2015 – July 2019
Post-Doctoral Researcher (until September 2015) / Research Staff Member *Yorktown Heights, NY, USA*

- CORAL Project
  – Worked on improving automatic vectorizer performance in LLVM compiler
  – Worked on designing XLFLANG, IR-to-IR translator for Fortran with full OpenMP 4.5 support
  – Worked on designing and evaluating efficient OpenMP code-generation scheme for nested parallelism on GPU
- PathForward Project
  – Working on improving performance of XLA, JIT compiler for Tensorflow
  – Working on designing neural networks trained with program IRs annotated with static and dynamic metadata
  – Working on designing runtime optimizations that leverage the neural networks trained with program IRs

University of Illinois, Urbana-Champaign January 2015 – May 2015
Post-Doctoral Researcher (Part-Time) *Urbana, IL, USA*

University of Illinois, Urbana-Champaign August 2008 – December 2014
Research Assistant *Urbana, IL, USA*

- DeNovo Project
  – Designed DeNovo, a shared-memory architecture that leverages disciplined programming models for better complexity- and energy-efficiency.
  – Developed an optimization for DeNovo cache coherence protocol, i.e., flexible-granularity data transfer using region information.
  – Designed and implemented DeNovoND, an extension for DeNovo that provides support for non-deterministic codes while maintaining the advantages of DeNovo.
  – Designed and implemented DeNovoSync, a hardware mechanism that provides support for racy synchronization accesses on DeNovo without writer-initiated invalidations.
- DPJ (Deterministic Parallel Java) Project
  – Ported and tested Java and non-Java benchmarks from Java Grande, SPLASH-2, and STAMP into DPJ.
  – Working on extending DPJ to support pipeline parallelism.
- ParKD (Parallel k-D tree construction) Project
  – Designed and implemented a parallel k-D tree construction algorithm exploiting nested parallelism by using Intel Threading Building Block (TBB) library.

Intel Corporation June 2010 – August 2010
Graduate Intern *Champaign, IL, USA*

- Wrote parallel TBB versions for Intel-internal "Parallel Language Futures" benchmarks.
- Conducted performance tuning for TBB, Cilk, OpenMP, CnC and Ct versions of the benchmarks, analyzing strengths and weaknesses of each run-time.

University of California, San Diego June 2006 – August 2008
Research Assistant *La Jolla, CA, USA*

- Designed and implemented a MATLAB® compiler front-end for tiled microprocessors.
- Extended MAGICA (shape and type inference engine for MATLAB®) by designing post-processing algorithms that provide more accurate inference results.
- Implemented loop induction variable detection and privatization algorithms for the back-end.

Samsung Electronics August 2004 – July 2006
Research Engineer *Suwon, Kyunggi-do, South Korea*

- ▷ Participated in the standardization process by Blu-ray® Disc Associations.
- ▷ Designed and developed a universal player for both Blu-ray® and HD-DVD®.
- ▷ Developed a DMA (Defect Management Area) Verifier, integrity-checking SW of Blu-ray disc at physical sector level.

RESEARCH INTERESTS

Deep learning for compilers and runtime
– Low-level auto-tuning for deep learning compiler frameworks
– DL-based performance predictor models
– Graph optimizations for DL models on large-scale CGRA accelerators

Compiler and runtime support for PIM/NDP hardware
– Runtime scheduling and placement mechanisms for PIM hardware
– PIM IR for specialized optimizations and code generation

Memory hierarchy for parallel systems
– Software-driven design of cache coherence protocols
– Signature based tracking and propagation for coherence information
– Hardware optimizations for handling synchronization accesses

Compiler technologies
– Compiler optimizations for high-level parallel constructs
– Code generation for offloaded device functions in high-level parallel languages
– Automatic vectorization
– IR-to-IR translation

Parallel programming
– Parallel algorithms that are performance portable across different coherence schemes

PUBLICATIONS

Jaehun Ryu, Eunhyeok Park, Hyojin Sung, *One-Shot Tuner for Deep Learning Compilers*, In *the 31st ACM SIGPLAN International Conference on Compiler Construction*, 2022.

Hyungkyu Ham, Hyunuk Cho, Minjae Kim, Jueon Park, Jeongmin Hong, Hyojin Sung, Eunhyeok Park, Euicheol Lim, Gwangsun Kim, *Near-Data Processing in Memory Expander for DNN Acceleration on GPUs*, In *IEEE Computer Architecture Letters*, Nov 2021.

Yongwon Shin, Hyojin Sung, *Hybrid Register Allocation with Spill Cost and Pattern Guided Optimization*, *In 34th International Workshop in Languages and Compilers for Parallel Computing (LCPC)*, 2021.

Hyojin Sung, Tong Chen, Alex Eichenberger, Kevin K. O'Brien, *CogR: Exploiting Program Structures for Machine-Learning Based Runtime Solutions*, In *the 28th International Conference on Parallel Architectures and Compilation Techniques (PACT)*, 2019 (poster).

A. C. Jacob, A. E. Eichenberger, H. Sung, S. F. Antao, G. Bercea, C. Bertolli, A. Bataev, T. Jin, T. Chen, Z. Sura, G. Rokos, K. O'Brien, *Efficient Fork-Join on GPUs Through Warp Specialization*, In *IEEE 24th International Conference on High Performance Computing (HiPC)*, 2017.

Hyojin Sung, Tong Chen, Zehra Sura, and Tarique Islam, *Leveraging OpenMP 4.5 Support in CLANG for Fortran*, In *the 13th international workshop on OpenMP (IWOMP)*, 2017.

Gheorghe-Teodor Bercea, Carlo Bertolli, Arpith C. Jacob, Alexandre E. Eichenberger, Alexey Bataev, Georgios Rokos, Hyojin Sung, Tong Chen, Kevin O'Brien, *Implementing implicit OpenMP data sharing on GPUs*, In *the Fourth Workshop on the LLVM Compiler Infrastructure in HPC at Super Computing (SC)*, 2017.

Samuel F. Antão, Alexey Bataev, Arpith C. Jacob, Gheorghe-Teodor Bercea, Alexandre E. Eichenberger, Georgios Rokos, Matt Martineau, Tian Jin, Guray Ozen, Zehra Sura, Tong Chen, Hyojin Sung, Carlo Bertolli,

Kevin O'Brien, *Offloading Support for OpenMP in Clang and LLVM*, In *the Fourth Workshop on the LLVM Compiler Infrastructure in HPC at Super Computing (SC)*, 2016.

Matt Martineau, Simon McIntosh-Smith, Carlo Bertolli, Arpith C. Jacob, Samuel F. Antão, Alexandre E. Eichenberger, Gheorghe-Teodor Bercea, Tong Chen, Tian Jin, Kevin O'Brien, Georgios Rokos, Hyojin Sung, Zehra Sura, *Performance Analysis and Optimization of Clang's OpenMP 4.5 GPU Support*, In *the 7th International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS16) at Super Computing (SC)*, 2016.

Tong Chen, Zehra Sura, and Hyojin Sung, *Automatic Copying of Pointer-based Data Structures*, In *the 29th International Workshop on Languages and Compilers for Parallel Computing (LCPC)*, 2016.

Carlo Bertolli, Samuel Antao, Gheorghe-Theodor Bercea, Arpith Jacob, Alexandre Eichenberger, Tong Chen, Zehra Sura, Hyojin Sung, Georgious Rokos, David Appelhans and Kevin O'Brien, *Integrating GPU Support for OpenMP Offloading Directives into Clang*, In *the Second Workshop on the LLVM Compiler Infrastructure in HPC at Super Computing (SC)*, 2015.

Gheorghe-Teodor Bercea, Carlo Bertolli, Samuel F. Antao, Arpith C. Jacob, Alexandre E. Eichenberger, Tong Chen, Zehra Sura, Hyojin Sung, Georgios Rokos, David Appelhans and Kevin O'Brien, *Performance Analysis of OpenMP on a GPU Using a CORAL Proxy Application*, In *the 6th International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS15) at Super Computing (SC)*, 2015.

Arpith C. Jacob, Samuel F. Antao, Hyojin Sung, Alexandre Eichenberger, Carlo Bertolli, Gheorghe-Teodor Bercea, Tong Chen, Zehra Sura, Georgios Rokos, and Kevin O'Brien, *Towards Performance Portable GPU Programming with RAJA*. In *the Workshop on Portability Among HPC Architectures for Scientific Applications at Super Computing (SC)*, 2015.

Robert Smolinski, Rakesh Komuravelli, Hyojin Sung, and Sarita V. Adve, *Eliminating On-Chip Traffic Waste: Are We There Yet?*, In *2015 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, 2015 (poster presentation).

Hyojin Sung and Sarita V. Adve, *Supporting Arbitrary Synchronization without Writer-Initiated Invalidations*, In *20th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2015.

Hyojin Sung, Rakesh Komuravelli, Sarita V. Adve, *DeNovoND: Efficient Hardware Support for Disciplined Non-Determinism*, In **IEEE Micro Top Picks from the Computer Architecture Conferences**, May/June 2014.

Hyojin Sung, Rakesh Komuravelli, Sarita V. Adve, *DeNovoND: Efficient Hardware Support for Disciplined Non-Determinism*, In *18th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2013.

Byn Choi, Rakesh Komuravelli, Hyojin Sung, Robert Smolinski, Nima Honarmand, Sarita V. Adve, Vikram S. Adve, Nicholas P. Carter, and Ching-Tsun Chou, *DeNovo: Rethinking the Memory Hierarchy for Disciplined Parallelism*, In *20th International Conference on Parallel Architectures and Compilation Techniques (PACT)*, 2011 (citation count: 72). **Best Paper Award.**

Byn Choi, Rakesh Komuravelli, Victor Lu, Hyojin Sung, Robert L. Bocchino, Sarita V. Adve, and John C. Hart, *Parallel SAH k-D Tree Construction*, In *High Performance Graphics (HPG)*, 2010 (citation count: 67).

Byn Choi, Rakesh Komuravelli, Hyojin Sung, Robert Bocchino, Sarita V. Adve, and Vikram V. Adve, *DeNovo: Rethinking Hardware for Disciplined Parallelism*, In *Second USENIX Workshop on Hot Topics in Parallelism (HotPar)*, 2010.

Robert Bocchino, Vikram Adve, Danny Dig, Sarita Adve, Stephen Heumann, Rakesh Komuravelli, Jeffrey Overbey, Patric Simmons, Hyojin Sung, and Mohsen Vakillian, *A Language for Deterministic-by-Default Parallel Programming*, In *15th Workshop on Compilers and Parallel Computing (CPC)*, 2010.

Robert Bocchino, Vikram Adve, Danny Dig, Sarita Adve, Stephen Heumann, Rakesh Komuravelli, Jeffrey Overbey, Patric Simmons, Hyojin Sung, and Mohsen Vakillian, *A Type and Effect System for Deterministic Parallel Java*, In the *Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, 2009 (citation count: 308).

Hyojin Sung, *A Portable MATLAB Front-End for Tiled Microprocessors*, Master Thesis, http://roger.ucsd.edu/record=b6636516~S7, 2009.

Alan Cooper, Translated into Korean by Hyojin Sung, *Inmates are running the asylum*, 2004.

PROFESSIONAL TALKS AND POSTERS

CogR: Exploiting Program Structures for Machine-Learning Based Runtime Solutions, poster presentation in PACT 2019.

Efficient Fork-Join on GPUs Through Warp Specialization, in HiPC 2017.

Leveraging OpenMP 4.5 Support in CLANG for Fortran, in IWOMP 2017.

Automatic Copying of Pointer-based Data Structures, at Programming Day at IBM 2016.

DeNovo: Rethinking the Memory Hierarchy for Disciplined Parallelism, candidate talk for AMD research, October 2014.

DeNovo: Rethinking the Memory Hierarchy for Disciplined Parallelism, candidate talk for IBM research, October 2014.

Rethinking Multicores for Disciplined Parallelism in the 2nd Rising Stars in EECS workshop at MIT, 2013.

DeNovoND: Efficient Hardware Support for Disciplined Non-Determinism in ASPLOS 2013.

DeNovo: Rethinking the Multicore Memory Hierarchy for Disciplined Parallelism in the best paper session of PACT 2011, October 2011.

DeNovo: Rethinking the Multicore Memory Hierarchy for Disciplined Parallelism in the "UPCRC Seminar" series, sponsored by Intel and Microsoft, December 2010.

DeNovo: Rethinking Hardware for Disciplined Parallelism in the "UPCRC Illinois Summit", sponsored by Intel and Microsoft, March 2010.

PATENTS

Alper Buyuktosunoglu and Hyojin Sung. *Anomalous Cache Coherence Transaction Detection in a Heterogeneous System*, filed in 2020.

Li Zhang, Chen Tong, Hyojin Sung, and Tian Jin. *Intelligent Space Exploration for Program Optimization with additional Compiler/Runtime Information*, filed in 2019.

Zehra Sura, Tong Chen, and Hyojin Sung. *Embedded program representation as input for machine learning applications*, filed in 2018.

Sung-hee Hwang, Hyojin Sung, Sung-ryeul Rhyu, and Jungwan Ko, *Verification Method and Apparatus*, US Patent US7617426, filed in 2006 and published in 2008.

SOFTWARE DISTRIBUTED

One-Shot Tuner for Deep Learning Compilers https://doi.org/10.5281/zenodo.6337971
Parallel k-D Tree Construction https://github.com/bchoi/ParKD

PROFESSIONAL SERVICE

Program committee member for LCTES 2022.

Student Research Competition (SRC) chair for CGO 2021.

Program committee member for ISPASS 2019.

External review committee member for HPCA 2018.

Reviewer for Transactions on Parallel and Distributed Systems (TPDS) 2017.

External review committee member for ISCA 2017.

Peer reviewer for IEEE Computer Architecture Letters in 2015.

Submission co-chair for ASPLOS 2014.

Student reviewer for OOPSLA 2009.

OTHER ACTIVITIES

PhD Mentor for the PhD mentoring program by the Department of Computer Science at UIUC in Fall 2012.

Student travel grant for HotPar 2010, ISCA 2011, PACT 2011, ASPLOS 2013, and ASPLOS 2014.

President of the Korean Buddhist Student Association at UIUC from August 2010 to July 2011.