

# Min Lee

12345 Alameda Trace Circle 816  
Austin, TX, 78727

512-529-8976  
min.lee@gatech.edu  
<http://osinside.net/mlee/>

---

## Education

- Ph.D., Computer Science, Georgia Institute of Technology, Atlanta, GA December 2013
- M.S., Computer Science, KAIST, Daejeon, Korea February 2006  
Best paper award of 2006 graduation, Computer Science Department, KAIST
- B.S., Computer Science, Yonsei University, Seoul, Korea February 2004  
Graduation with highest honor in Computer Science Department

## Research Areas / Summary

- Hypervisor, System virtualization – Xen, VM scheduling and consolidation, Memory and cache management, MMU virtualization, Real-time support
- Operating System – Linux, Physical memory management, Buffer cache
- Computer Architecture – Multi/manycore, last-level/distributed cache management
- Cache and memory hierarchy – NUMA memory support, interconnect
- Parallel computing and runtime – HPC workload support at runtime/OS, MPI library
- Emerging memory technologies – Software support for new memory technologies (such as 3D-stacked)
- Power-aware system – Power-aware memory support

## Papers

- Vishal Gupta, Min Lee, Karsten Schwan. "HeteroVisor: Exploiting Resource Heterogeneity to Enhance the Elasticity of Cloud Platforms." VEE 2015, Istanbul, Turkey, March 14-15, 2015.
- Joong-Yeon Cho, Hyun-Wook Jin, Min Lee and Karsten Schwan. "Dynamic Core Affinity for High-Performance File Upload on Hadoop Distributed File System." Parallel Computing, the special issue on Data-Intensive Scalable Computing Systems, 2014.
- Joong-Yeon Cho, Hyun-Wook Jin, Min Lee and Karsten Schwan. "On the Core Affinity and File Upload Performance of Hadoop." DISCS 2013 (The 2013 International Workshop on Data-Intensive Scalable Computing Systems), Denver, CO, November 18, 2013.
- Min Lee, Vishal Gupta, Karsten Schwan. "Software-Controlled Transparent Management of Heterogeneous Memory Resources in Virtualized Systems." MSPC 2013 (ACM SIGPLAN Workshop on Memory Systems Performance and Correctness), Seattle, Washington, June 21, 2013.
- Min Lee, Karsten Schwan. "Region Scheduling: Efficiently Using the Cache Architectures via Page-level Affinity." ASPLOS 2012, London, UK, March 3-7, 2012.
  - ACM SIGARCH Computer Architecture News 40 (1), 451-462
  - ASPLOS 2012 Student Travel Grant
- Mrinmoy Ghosh, Ripal Nathuji, Min Lee, Karsten Schwan, and Hsien-Hsin S. Lee. "Symbiotic Scheduling for Shared Caches in Multi-Core Systems Using Memory Footprint Signature." ICPP 2011 (IEEE International Conference on Parallel Processing), Taipei, Taiwan, September, 2011.
- Min Lee, A. S. Krishnakumar, P. Krishnan, Navjot Singh, Shalini Yajnik. "Hypervisor-Assisted Application Checkpointing in Virtualized Environments." DSN 2011 (International Conference on Dependable Systems and Networks), Hong Kong, June 27-30, 2011.
  - DSN 2011 Student Travel Grant
- Min Lee, A. S. Krishnakumar, P. Krishnan, Navjot Singh, Shalini Yajnik. "XenTune: Detecting Xen Scheduling Bottlenecks for Media Applications." IEEE Globecom 2010 (Communications Software, Services and Multimedia Applications Symposium), Miami, FL, Dec 6-10, 2010.

- Min Lee, A. S. Krishnakumar, P. Krishnan, Navjot Singh, Shalini Yajnik. "Supporting Soft Real-Time Tasks in the Xen Hypervisor." VEE 2010, Pittsburgh, PA, March 17-19, 2010.
  - ACM Sigplan Notices 45 (7), 97-108
  - Invited and presented at Xen Summit North America at AMD April 28-29, 2010.
- Jiantao Kong, Karsten Schwan, Min Lee, and Mustaque Ahamad. "ProtectIT: Trusted Distributed Services Operating on Sensitive Data." EuroSys 2008, Glasgow, Scotland, March 31 – April 4, 2008.
  - ACM SIGOPS Operating Systems Review 42 (4), 137-147
- Min Lee, Euiseong Seo, Joonwon Lee, Jin-soo Kim. "PABC:Power Aware Buffer Cache Management for Low Power Consumption." IEEE Transactions on computers, April 2007 (Vol. 56, No. 4) pp. 488-501.
  - Three awards received. (See KAIST bullet in Awards section.)

## Book

- Operating System Inside (In Korean)
  - <http://osinside.net/osinside/osinside.htm>

## Research/Work Experience

- Intel, Santa Clara, California 2013~Present
  - Memory management in Open Collaborative Runtime (OCR, see <https://xstack.exascale-tech.com/wiki/>) for Xstack project which aims novel task-based runtime for exascale HPC. This supports Intel's open-source research prototype TG architecture.
  - Design/Implement runtime/OS for Exascale project (FastForward), which aims entirely new high performance, energy-efficient next-generation processor that is essential for developing Exascale systems.
  - Design/Implement lightweight MPI library for the new novel Exascale processor to support legacy code.
- Avaya Labs, Basking Ridge, New Jersey (was part of Bell Labs) 2009/2010 Summer
  - Design/Implement Xen real-time scheduler [VEE2010] to support enterprise IP Telephony server workload including Xen monitoring tools [Globecom2010] and hypervisor-assisted application checkpointing work for high availability [DSN2011].
- Georgia Institute of Technology, Atlanta, Georgia 2006~2013
  - A study on core affinity for Hadoop [DISCS2013, PC2014]
  - Design/Implement SW management of heterogeneous memory (3D-stacked memory) in virtualized environment [MSPC2013, VEE2015].
  - Design/Implement software approach at hypervisor-level (Xen) for cache management in CMP environment. Departing from traditional CPU-centric scheduler, newly memory-centric scheduler is proposed [ASPLOS2012].
  - Design and evaluate shared-memory based communication among Xen-domains. Develop and evaluate kernel module for efficient communication [EuroSys2008].
- Intel, Portland, Oregon 2007 Summer
  - Design and analyze on NUCA-aware scheduling policy using SESC microprocessor simulator and Bochs full system simulator. To make efficient use of NUCA architecture for CMP processors, new OS policies are proposed. [ICPP2011]
- KAIST, Daejeon, Korea 2004~2006
  - Design/implement PABC (Power-Aware Buffer Cache) management scheme in Linux kernel [TOC2007], which led to my master thesis.
- ETRI (Electronics and Telecommunications Research Institute), Daejeon, Korea 2005 Fall
  - RFID/USN project hosted by ETRI
  - Implementing RFID reader protocols and tag processing engines
- Yonsei University, Seoul, Korea 1998~2004
  - Internet-based radio project for embedded devices. Personal Audio Station, Undergraduate thesis. Kernel scheduler and memory allocator modification and evaluation. First place prize in the Yonsei creativity fair.
- A co, 122Sig, 2<sup>nd</sup> Infantry Division, US 8<sup>th</sup> Army in Korea (Military service) 2000~2002

## **Fellowships, Grants, and Awards**

- Intel
  - Q2'2016 TCAR Award, TCAR's highest recognition, for critical OCR and tools support on XStack2
  - Q2'2015 DCG Award for a disruptive execution model implemented as a runtime system (OCR) helps the community to continue innovations, strengthening future Intel platforms.
  - Q2'2013 TCAR Award, the department's highest recognition, for that "ramp & delivery went above and beyond all expectations".
- KAIST
  - Best paper award of 2006 graduation, Computer Science Department, KAIST.
  - Bronze Prize in the Samsung Humantech Thesis Award competition, Feb 2006.
  - Prize in the 25th student's thesis competition, Korea Information Science Society. (2006)
- KFAS (Korea Foundation for Advanced Study)
  - Doctoral Study Abroad Scholarship (2004)
- Yonsei University
  - Scholarships for academic excellence (1999 Fall, 2002/2003 Spring)
  - Prizes with high honor (1999 Fall, 2003 Spring)
  - Prize with highest honor (2002 Spring)
  - Graduation with highest honor in CS department. (2003 Fall)
  - First place prize in the Yonsei creativity fair (Undergraduate thesis competition)

## **Skills/Experiences**

- Linux / Xen kernel development, MPI programming
- C/C++/Java/Python, and some LLVM/Clang experience
- A novel task-based parallel runtime, OCR, for high performance computing
- Newlib experience (C library for embedded systems)
- Rose compiler (source-to-source translator)

## **Software**

- Soft real-time scheduler for the Xen hypervisor
- MPI-lite library for next-generation chips for HPC workloads
- Quick Allocator for fast and scalable memory allocation

## **Patents**

- U.S. Patent #8,245,234 Credit scheduler for ordering the execution of tasks
- U.S. Patent #8,166,485 Dynamic techniques for optimizing soft real-time task performance in virtual machines
- U.S. Patent #8,161,491 Soft real-time load balancer

## **Extra Activities**

- Organizing/leading Linux kernel study groups in Yonsei CS dept.
- KTF mobile futurist activity (internship)
- Winner of KTF-cup Alkkagi Tournament.
- Small College 20th class. (Classic reading and discussion activity)
- President of GTKSA-CS (GT Korean Students Association in Computer Science) in 2007-2008

*Last Updated: June 2016*