Gunjae Koo

CONTACT INFORMATION

Address:
114 Gangchon-ro Apt. 512-1406
Goyang-si, Gyeonggi-do 10419, South Korea

Mobile: +82-10-3896-0708
E-mail: gunjae.koo@hongik.ac.kr
Homepage: http://gunjaekoo.com

EDUCATION

University of Southern California, Los Angeles, CA

Ph.D., Electrical Engineering, Aug 2018

- Thesis: *Architectural Innovations for Mitigating Data Movement Cost on Graphics Processing units and Storage Systems*
- Adviser: Professor Murali Annavaram
- Area of Study: computer architecture, embedded systems, storage systems, near-storage data analytics, parallel processor architecture, high-performance computing

Seoul National University, Seoul, South Korea

M.S., Electrical Engineering and Computer Science, Feb 2003

- Thesis: *An Equalizer and Viterbi Decoder Design for 1000BASE-T Gigabit Ethernet*
- Adviser: Professor Deog-Kyoon Jeong
- Area of Study: digital systems design, VLSI, model-based design, signal processing, communication theory

B.S., Electrical Engineering, Feb 2001

- Thesis: *2.4GHz Low Noise Amplifier Design in a 0.25 μm CMOS Technology*
- Electrical specialization (emphasis on digital CMOS circuits design, digital systems design, computer systems and signal processing)

PROFESSIONAL EXPERIENCE

Hongik University, Seoul, South Korea

Assistant Professor (Tenure Track) **Sep 2018 to present**

- Area of Research: computer architecture, memory & storage systems, embedded systems, FPGA

University of Southern California, Los Angeles, CA

Graduate Research Assistant **Aug 2012 to Aug 2018**

- Advisor: Professor Murali Annavaram
- Area of Research: computer architecture, embedded systems, storage systems, near data processing, GPGPU, memory systems, energy efficient computing

Intel, Hillsboro, OR

Graduate Research Intern **May 2016 to Dec 2016**

- Area of Research: memory controller for server architecture, next-generation memory technology

LG Electronics, Seoul, South Korea

Senior Research Engineer **Apr 2008 to May 2011**

- SoC Group. System IC Laboratory / Digital TV Laboratory
- Area of Research: SoC design, FPGA platform design, digital systems architecture & modeling, RTL design, verification & validation, image processing algorithm, image analysis

Senior Research Engineer **Jul 2005 to Apr 2008**

- SoC Core Technology Group, Device & Material Laboratory
- Area of Research: SoC design, digital systems architecture & modeling, RTL design, signal processing algorithm

Junior Research Engineer **Mar 2003 to Jul 2005**

- Digital Media ASIC group, Digital Storage Research Laboratory
- Area of Research: signal processing algorithm, communication theory, RTL design, verification

Seoul National University, Seoul, South Korea

Graduate Research Assistant **Mar 2001 to Feb 2003**

- Advisor: Professor Deog-Kyoon Jeong
- Area of Research: digital systems design, verification, testing, signal processing algorithm, communication theory

Bitnuri, Seoul, South Korea

Researcher (co-founder) **Mar 2000 to Dec 2000**

- Area of Research: digital systems design

PUBLICATIONS

Refereed Conferences

[1] Gunjae Koo, Hyeran Jeon, Zhenhong Liu, Nam Sung Kim, and Murali Annavaram. "CTA-Aware Prefetching and Scheduling for GPU". *Proceedings of the 32nd IEEE International Parallel and Distributed Processing Symposium (IPDPS '18)*, May 21–25, 2018.

[2] Gunjae Koo, Kiran Kumar Matam, Te I, Hema Venkata Krishna Giri Nara, Jing Li, Hung-Wei Tseng, Steven Swanson and Murali Annavaram. "Summarizer: Trading Communication with Computing Near Storage". *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO '17)*, Oct 14–18, 2017.

[3] Gunjae Koo, Yunho Oh, Won Woo Ro and Murali Annavaram. "Access Pattern-Aware Cache Management for Improving Data Utilization in GPU". *Proceedings of the 44th Annual International Symposium on Computer Architecture (ISCA '17)*, Jun 24–28, 2017.

[4] Keunsoo Kim, Sangpil Lee, Myung Kuk Yoon, Gunjae Koo, Won Woo Ro and Murali Annavaram. "Warped-Preexecution: A GPU Pre-execution Approach for Improving Latency Hiding". *Proceedings of the 22nd International Symposium on High Performance Computer Architecture (HPCA '16)*, Mar 12–16, 2016.

[5] Gunjae Koo, Hyeran Jeon and Murali Annavaram. "Revealing Critical Loads and Hidden Data Locality in GPGPU Applications". *Proceedings of the 2015 IEEE International Symposium on Workload Characterization (IISWC '15)*, Oct 4–6, 2015.

[6] Sangpil Lee, Keunsoo Kim, Gunjae Koo, Hyeran Jeon, Won Woo Ro and Murali Annavaram. "Warped-Compression: Enabling Power Efficient GPUs through Register Compression". *Proceedings of the 42nd Annual International Symposium on Computer Architecture (ISCA '15)*, Jun 13–17, 2015.

[7] Gunjae Koo, Kyoung Won Lim and Seung Jong Choi. "Complementary Block-Based Motion Estimation for Frame Rate Up-Conversion". *The 2011 IEEE International Conference on Consumer Electronics (ICCE '11)*, Jan 9–12, 2011.

[8] Gunjae Koo, Woochul Jung and Heesub Lee. "A Robust PRML Read Channel with Digital Timing Recovery for Multi-Format Optical Disc". *Proceedings of the 2006 IEEE International Symposium on Circuits and Systems (ISCAS '06)*, May 21–24, 2006.

[9] Gunjae Koo, Woochul Jung and Heesub Lee. "A PRML Read Channel for High Density Optical Disc". *The 13th Korean Conference on Semiconductors*. Feb 23–24, 2006.

Journals

[10] Sangpil Lee, Keunsoo Kim, Gunjae Koo, Hyeran Jeon, Murali Annavaram and Won Woo Ro. "Improving Energy Efficiency of GPUs through Data Compression and Compressed Execution". *IEEE Transactions on Computers*, Vol. 66, No. 5, pp. 834–847. 2017.

[11] Kyoung Won Lim, Hansoo Kim, Hyunchul Noh, Hyunchul Shin, Woochul Jung, Gunjae Koo and Ryuk Park. "Advanced Frame Rate Conversion without Halo and Judder Effect for 120Hz LCD Displays". *iMiD/IDMC/Asia Display*, Vol. 8 Book-II, pp. 1397–1400. 2008.

Workshops and Technical Reports

[12] Gunjae Koo, Kiran Kumar Matam, Te I, Hema Venkata Krishna Giri Nara, Jing Li, Hung-Wei Tseng, Steven Swanson, and Murali Annavaram. "Dynamic Near Data Processing Framework for SSDs". *The 9th Annual Non-Volatile Memories Workshop (NVMW '18)*, Mar 11–13, 2018.

[13] Hyeran Jeon, Gunjae Koo and Murali Annavaram. "CTA-aware Prefetching for GPGPU". *Computer Engineering Technical Report, Ming Hsieh Department of Electrical Engineering, University of Southern California*, Oct, 2014.

Thesis

[14] Gunjae Koo. "Architectural Innovations for Mitigating Data Movement Cost on Graphics Processing units and Storage Systems". *Ph.D. Dissertation. Department of Electrical Engineering, University of Southern California*, Aug, 2018.

[15] Gunjae Koo. "An Equalizer and Viterbi Decoder Design for 1000BASE-T Gigabit Ethernet". *M.S. Thesis, Department of Electrical Engineering and Computer Science, Seoul National University*, Feb, 2003.

PATENTS

U.S. Patents

[1] Gunjae Koo, Vivek Kozhikkottu, Shankar Ganesh Ramasubramanian, Christopher B. Wilkerson. "Increasing Read Pending Queue Capacity to Increase Memory Bandwidth". *U.S. Patent Pub. No. US 2018/0188976 A1*, Jul 5, 2018.

[2] Gunjae Koo and Eun Pyo Lee. "Frequency Detection Method for Optical Disc Bit Reproduction Apparatus". *U.S. Patent No. US 7,433,289*, Oct 7, 2008.

Korean Patents

[3] Gunjae Koo. "Motion Estimation Method for Video Signal". *Korea Patent 10-2012-0106279*, Sep 26, 2012.

[4] Gunjae Koo. "Method of Frame Interpolation by Complementary Motion Estimation Algorithm". *Korea Patent 10-2011-0034241*, Apr 5, 2011.

[5] Gunjae Koo. "Method of Assigning Motion Vector of Occlusion Region". *Korea Patent 10-2011-0034242*, Apr 5, 2011.

[6] Gunjae Koo and Heesub Lee. "Liquid Crystal Display". *Korea Patent 10-2007-0077746*, Jul 27, 2007.

[7] Gunjae Koo and Hansoo Kim. "Method for Generating CAV Clock of Optical Disc". *Korea Patent 10-2007-0028754*, Mar 13, 2007.

[8] Gunjae Koo and Hyugjin Kwon. "Adaptive Boost Gain Controlled Limit Equalizer and Gain Calculation Method in Above Limit Equalizer". *Korea Patent 10-2006-0073662*, Oct 12, 2006.

[9] Gunjae Koo. "Optical Recording System for Detecting Frequency of Read Channel Using Wobble Signal and Playback Method Thereof". *Korea Patent 10-2006-0036534*, Jun 15, 2006.

[10] Gunjae Koo and Eun Pyo Lee. "Frequency Detecting Method in Optical Disk Bit Data Reproducing System". *Korea Patent 10-2005-0080866*, Jun 23, 2006.

[11] Gunjae Koo and Youngsoo Jang. "Apparatus of Detecting in Optical Disc and Method of Same". *Korea Patent 10-2005-0076072*, Apr 28, 2006.

TALKS

[1] Diving into Data: Architectural Approaches to In-Storage Computing.
KIISE Computer System Society Conference, Pyeongchang, South Korea, Jan 2019.

[2] GPU Memory System Architecture for Big Data.
EDA Winter Workshop, Pyeongchang, South Korea, Jan 2019.

[3] CTA-Aware Prefetching and Scheduling.
IPDPS '18, Vancouver, British Columbia, Canada, May 2018.

[4] Dynamic Near Data Processing Framework for SSDs.
NVMW '18, La Jolla, California, Mar 2018.

[5] Architectural Challenges and Innovation for Accelerating Big Data Analytics.
Yonsei University, Seoul, South Korea, Sep 2018.
Postech, Pohang, South Korea, Jun 2018.
Texas A&M University, College Station, Texas, Apr 2018.
University of Central Florida, Orlando, Florida, Apr 2018.
Binghamton University, Binghamton, New York, Mar 2018.
University of California Santa Cruz, Santa Cruz, California, Mar 2018.

[6] Summarizer: Trading Communication with Computing Near Storage.
MICRO '17, Cambridge, Massachusetts, Oct 2017.

[7] Access Pattern-Aware Cache Management for Improving Data Utilization in GPU.
ISCA '17, Toronto, Ontario, Canada, Jun 2017.

[8] Revealing Critical Loads and Hidden Data Locality in GPGPU Applications.
IISWC '15, Atlanta, Georgia, Oct 2015.

[9] Complementary Block-Based Motion Estimation for Frame Rate Up-Conversion.
ICCE '11, Las Vegas, Nevada, Jan 2011.

RESEARCH PROJECTS

University of Southern California, Los Angeles, CA

Smart Storage Systems **Jan 2017 to present**

Data movement cost from storage devices to compute nodes is extremely high for modern data processing applications. Such big data applications pay a significant fraction of their execution time on the data input/output (I/O) time. Near data processing (NDP) is a prominent approach by offloading part of computation to embedded storage processors. NDP is becoming a more viable option with SSDs, which equip embedded processors and page buffer memory. This computation potential in the embedded storage processors can facilitate *active storage systems*, which enable computing near storage or application-specific storage data management.

- *Summarizer – a dynamic NDP framework (**MICRO** '17)*

GPU Memory Hierarchy Architecture **Sep 2014 to May 2016**

GPU's memory operation is a critical performance bottleneck since lots of memory requests are issued from dozens of warps (or wavefronts) within a short time window. Thus, limited resources in memory subsystem suffer from low efficiency due to severe data contention and heavy data traffic. I characterized diverged memory requests in GPGPU applications to reveal main performance bottlenecks from the critical load instructions. I also investigated GPU's unique data access patterns by global loads. Based on these observations I presented the GPU cache management method to utilize the data cache more efficiently.

- *Access Pattern-Aware Cache Management (**ISCA** '17)*
- *Revealing Critical Loads (**IISWC** '15)*
- *Warped-Compression (**ISCA** '15)*

Prefetch **Sep 2013 to May 2015**

Long latency of memory operation is one of critical performance hurdles in general computing processors including GPUs. Prefetch can be one of prominent approach to hide this long latency of data fetch. I proposed an efficient GPU prefetch mechanism based on GPU's unique software execution model and array index calculation approaches. I also suggested a warp scheduling scheme that can enhance the timeliness of the prefetcher.

- *CTA-Aware Prefetching and Scheduling (**IPDPS** '18)*
- *Warped-Preexecution (**HPCA** '16)*

Energy Efficient Computing **Apr 2012 to Aug 2013**

Battery efficiency varies with usage patterns of mobile computing due to the non-ideal characteristics in converting chemical to electrical energy. Besides the efficiency of power delivery network decides actual lifetime of batteries. I measured end-point power usage patterns and battery discharge time by mobile applications. I presented the control method applying DVFS to CPUs, GPUs, and peripherals to maximize the efficiency of battery and power delivery network.

Intel, Hillsboro, OR

Memory Controller for Server Architecture **May 2016 to Dec 2016**

I investigated the memory controller architecture for high-performance server processors and next-generation DRAM. Based on the performance analysis using a wide range of datacenter server application traces, I figured out performance issues that lead to high data access latency. I proposed several solutions that can alleviate read/write interventions to lower DRAM access latency.

LG Electronics, Seoul, South Korea

Frame Rate Conversion SoC **Jul 2006 to May 2011**

I developed the hardware architecture and RTL design of the enhanced motion estimation processor (eMEP) employed in the frame rate conversion (FRC) SoC. The *3rd generation* of FRC SoC included more function such as 2D-to-3D image conversion, 3D depth control and LED local dimming. I revised my motion estimation algorithm to support 3D image frames for the latest generation SoC. I made a hardware equivalent model of the motion estimation algorithm to implement eMEP hardware design. I verified the implemented hardware blocks with the in-house FPGA platform and ZeBu ASIC emulation system. To the best of my knowledge, the *1st generation* FRC SoC is the world's first one chip solution for real 200/240Hz FRC. FRC SoCs were deployed in the high-end 200/240Hz 3D TV sets, which were awarded in 2010 CES.

Digital TV main SoC **May 2009 to Dec 2009**

While at Digital TV Laboratory I participated in the DTV main SoC research project. The purpose of the project was developing a SoC that integrates most of the DTV functions such as DTV signal processing, image processing, and frame control within one chip package. I developed the hardware architecture and RTL design of the motion estimation processor (MEP). I verified the implemented design with the custom FPGA platform and ZeBu ASIC emulation system. I also fulfilled ASIC back-end tasks such as timing closure and gate-level debugging.

Frame Rate Up-conversion Algorithm **Jul 2005 to Apr 2008**

While at Digital TV Laboratory I developed a cost-effective motion estimation algorithm for frame rate up-conversion, which can reduce judder and halo effects without artifacts on the flat panel digital TV sets. I proposed an efficient block-based complementary motion estimation algorithm, which enables accurate motion detection between frames without heavy computation.

- *Complementary Motion Estimation (**ICCE** '11)*

High-Density and Multi-Format Optical Disc Player and Recorder Front-End SoCs
Mar 2003 to Jul 2005

While at Digital Storage Research Laboratory I developed hardware and algorithms for the data read channel employed in the front-end SoC for optical disk playback and recording. I proposed signal processing algorithms for modulator, demodulator, digital timing recovery, equalizer, and PRML decoder supporting multi-format optical discs (CD, DVD, and Blu-ray disc). I also designed the hardware architecture and RTL codes base on the proposed signal processing algorithms. I verified the hardware design with the custom FPGA platform to demonstrate the stable playback of optical discs that have damages and defects.

- *PRML Read Channel and Digital Timing Recovery (**ISCAS** '06)*

Seoul National University, Seoul, South Korea

Gigabit Ethernet Physical Layer Chip **Jul 2002 to Feb 2003**

I proposed digital equalizer, Viterbi decoder, and timing recovery algorithms that meet IEEE 802.3ab specification supporting Ethernet PHY for 10/100/1000 BASE-T using UTP-5 cables. The proposed design was verified with a wide range of channel parameters. I implemented the digital hardware blocks that include the proposed digital processing algorithms and the PHY controller layer for Gigabit Ethernet controller chip. I presented the proposed digital equalizer and Viterbi decoder design as my thesis for Master's degree.

Bitnuri, Seoul, South Korea

Flash ROM Controller for New Music Contents Media **Mar 2000 to Dec 2000**

I was one of the members leading the project awarded by *Korea University Students Venture Contest*. I designed the flash ROM controller employed to music contents media, which can substitute conventional optical storage media.

TEACHING EXPERIENCE

Hongik University, Seoul, South Korea

Assistant Professor

- 106824: Storage System Architecture (IT System Design) Spring 2019
- 106202: Digital Logic Design Spring 2019
- 106601: Computer Systems Organization Fall 2018
- 106610: Computer Communications Network Fall 2018
- 106818: Embedded Systems Design Fall 2018

University of Southern California, Los Angeles, CA

Teaching Assistant

- EE 354: Introduction to Digital Circuits — Fall 2015, Spring 2016
  - Responsible for lecture and supervision of laboratory
  - Guiding term projects to implement digital designs on FPGA boards
- EE 560: Digital System Design - Tools and Techniques — Summer 2013
  - Responsible for lecture and supervision of laboratory
  - Design and implementation of the simple MOESI cache coherence protocol for course material
- EE 557: Computer Systems Architecture — Fall 2013, Spring 2014
  - Responsible for 1-hour discussion and grading
  - Design of assignments and class projects

Students Mentored

- Rohit Madan, University of Southern California (Direct Research, Fall 2017)
- Yung-Hung Chen, University of Southern California (Direct Research, Fall 2017)
- Qili Wang, University of Southern California (Intern, Summer 2017)
- Chirag Ahuja, University of Southern California (Direct Research, Spring 2015)
- Kevin Jia, University of Southern California (Intern, Summer 2014)
- Sangmin Kim, University of Southern California (Direct Research, Spring 2014)

Inter-university Semiconductor Research Center, Seoul, South Korea

Teaching Assistant

- VLSI Design Coursework — Jan 2002

Seoul National Universit, Seoul, South Korea

Teaching Assistant

- Electronics Circuits — Spring 2001
  - Responsible for 1-hour discussion and grading of assignments
- Analog System Laboratory — Fall 2001
  - Responsible for supervision of 3-hour laboratory and grading of reports

PROFESSIONAL SERVICES

Conference/Workshop Session Chair

- 9th Annual Non-Volatile Memories Workshop (*NVMW '18*), Mar 2018.

Panelist

- USC Viterbi PhD Academic Career Mentoring Panel, May 2018.

Reviewer/External Reviewer

- International Symposium on Computer Architecture (*ISCA*)
- IEEE/ACM International Symposium on Microarchitecture (*MICRO*)
- ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems (*SIGMETRICS*)
- IEEE International Parallel and Distributed Processing Symposium (*IPDPS*)
- IEEE International Symposium on Workload Characterization (*IISWC*)
- IEEE International Conference on Computer Design (*ICCD*)
- IEEE International Symposium on Performance Analysis of Systems and Software (*ISPASS*)
- IEEE Computer Architecture Letters (*CAL*)
- ACM Transactions on Embedded Computing Systems (*TECS*)

AWARDS AND HONORS

University of Southern California

- Viterbi School of Engineering Doctoral Fellowship, Aug 2011–Jul 2015

LG Electronics

- The 2009 LG Group Best R&D Products Award: 4th prize, Mar 2010
- The 2008 Outstanding Researcher Award in DTV Research Laboratory, Jan 2009

Seoul National University

- Honors Scholarship in the Department of Electrical Engineering, Mar 1999

HARDWARE AND SOFTWARE SKILLS

Storage System Development Platforms

- Dragonfire board (Intelligent-SSD)
- OpenSSD

Computer Architecture Simulators

- Processor simulators: GPGPU-sim, DRAMSim, Ramulator, ZSim, Gem5, SimpleScalar, SESC, Multi2Sim.
- Power estimation: McPat, Wattch, CACTI.

SoC and VLSI Design

- Hardware description languages: Verilog, VHDL
- RTL simulation and verification tools: NCsim, Verdi, ModelSim
- RTL coverage and lint tools: ICCR, Spyglass
- Logic synthesis tool: Synopsys design compiler
- FPGA tools: Xilinx Vivado, Intel Quartus Prime, Chipscope
- ASIC emulation tool: ZeBu

Modeling and Analysis

- Analysis tool: Matlab
- Model-based design tool: Simulink

Programming

- Programming languages: C, C++, Python
- Script languages: bash shell, GNU make, Perl
- Version control tools: git, SVN, CVS

MEMBERSHIPS

ACM SIGARCH member
IEEE member
SID member
Samsung frontier membership
AYSO (American Youth Soccer Organization) referee

REFERENCES

References are available upon request.