Contact Information

School of Computer Science
Institute for Research in Fundamental Sciences (IPM)
Farmaniyeh Avenue
Tehran, Iran

http://cs.ipm.ac.ir/~plotfi
plotfi@ipm.ir
Tel: +98-21-24509404
Office: 401

Research Interests

Computer architecture. More specifically, I am interested in cross-stack and technology-driven innovations for improving performance and energy efficiency of computer systems for big-data applications.

Employment

- **Head of School of Computer Science.** Institute for Research in Fundamental Sciences (IPM) Apr. 2020
- **Associate Professor.** Institute for Research in Fundamental Sciences (IPM) 2019
- **Senior Postdoc Researcher.** Institute for Research in Fundamental Sciences (IPM) Feb. 2017
- **Director of Turin Cloud Services.** Institute for Research in Fundamental Sciences (IPM) Dec. 2016
- **Postdoc Researcher.** Institute for Research in Fundamental Sciences (IPM) Jul. 2014

Education

- **Ph.D. in Computer Science.** Swiss Federal Institute of Technology in Lausanne (EPFL) Sep. 2013
  Advisor: Prof. Babak Falsafi
  Thesis: *Scale-Out Processors*
- **M.S. in Electrical and Computer Engineering.** University of Tehran Feb. 2005
  Advisor: Prof. Zainalabedin Navabi
  Co-Advisor: Dr. Mehran Massoumi
  Thesis: *An Efficient Data Structure for RTL Representation*
  GPA: 18.76/20.00 (graduated with the highest honors)
- **B.S. in Electrical and Computer Engineering.** University of Tehran Sep. 2002
  Advisor: Prof. Zainalabedin Navabi
  Thesis: *Implementation of a VHDL-AMS-to-CHIRE Compiler*
  GPA: 17.40/20.00 (graduated with the highest honors)

Honors and Awards

- **Best Paper Award** from the 19^{th} International Symposium on Computer Architecture & Digital Systems (CADS) for "NOC Characteristics of Cloud Applications" 2017
- **Young Faculty Award** from Iran's National Elites Foundation 2016
- **Intel Ph.D. Fellowship Award.** 2012-2013 Academic Year 2012
- **Paper Award** from the European Network of Excellence on High Performance and Embedded Architecture and Compilation (HiPEAC) for "NOC-Out: Microarchitecting a Scale-Out Processor" 2012
- **Paper Award** from the European Network of Excellence on High Performance and Embedded Architecture and Compilation (HiPEAC) for "Scale-Out Processors" 2012
- **Paper Award** from the Department of Computer and Communication Sciences at EPFL for "Cuckoo Directory: A Scalable Directory for Many-Core Systems" 2011
- **Paper Award** from the European Network of Excellence on High Performance and Embedded Architecture and Compilation (HiPEAC) for "Cuckoo Directory: A Scalable Directory for Many-Core Systems" 2011
- **Best Student Paper Finalist** at the 17^{th} International Symposium on High-Performance Computer Architecture (HPCA) for "Cuckoo Directory: A Scalable Directory for Many-Core Systems" 2011
- **Dean's Honored Graduate.** College of Engineering, University of Tehran 2005
  Ranked 1^{st} among all computer-architecture students who graduated in 2005

- **Best Student Paper Award** from the 13^{th} Iranian Conference on Electrical Engineering for "Improving Logic-Level Representation of BMD/TED Diagrams" 2005
- **Ranked 8^{th}** among 5,445 participants of Iran's national M.S. entrance exam. Computer-Engineering track 2002
- **Dean's Honored Graduate.** College of Engineering, University of Tehran 2002
  Ranked 2^{nd} among all computer-engineering students who graduated in 2002
- **Faculty of Engineering (FOE) Award.** College of Engineering, University of Tehran 2001
  Ranked 1^{st} in computer-engineering track
  Annually awarded to the top three students in each track by College of Engineering
- **Ranked 11^{th}** among 30,000 participants of Iran's Region 3 B.S. entrance exam. Mathematics-and-Physics track 1998

PUBLICATIONS

Conference Papers

1. A. Ansari, **P. Lotfi-Kamran**, and H. Sarbazi-Azad, "Divide and Conquer Frontend Bottleneck," in *International Symposium on Computer Architecture (ISCA)*, pp. 65–78, June 2020.
2. M. Bakhshalipour, M. Shakerinava, **P. Lotfi-Kamran**, and H. Sarbazi-Azad, "Bingo Spatial Data Prefetcher," in *International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 399–411, February 2019.
3. A. Yazdanbakhsh, C. Song, J. Sacks, **P. Lotfi-Kamran**, N. S. Kim, and H. Esmaeilzadeh, "In-DRAM Near-Data Approximate Acceleration for GPUs," in *International Conference on Parallel Architectures and Compilation Techniques (PACT)*, pp. 34:1–34:14, November 2018.
4. M. Bakhshalipour, **P. Lotfi-Kamran**, and H. Sarbazi-Azad, "Domino Temporal Data Prefetcher," in *International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 131–142, February 2018.
5. **P. Lotfi-Kamran**, M. Modarressi, and H. Sarbazi-Azad, "NOC Characteristics of Cloud Applications," in *International Symposium on Computer Architecture & Digital Systems (CADS)*, pp. 18–23, December 2017. (**Selected by the program committee as the best paper of the conference**)
6. **P. Lotfi-Kamran**, M. Modarressi, and H. Sarbazi-Azad, "Near-Ideal Networks-on-Chip for Servers," in *International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 277–288, February 2017.
7. A. Yazdanbakhsh, J. Park, H. Sharma, **P. Lotfi-Kamran**, and H. Esmaeilzadeh, "Neural Acceleration for GPU Throughput Processors," in *International Symposium on Microarchitecture (MICRO)*, pp. 482–493, December 2015.
8. **P. Lotfi-Kamran**, B. Grot, and B. Falsafi, "NOC-Out: Microarchitecting a Scale-Out Processor," in *International Symposium on Microarchitecture (MICRO)*, pp. 177–187, December 2012.
9. D. Milojevic, S. Idgunji, D. Jevdjic, E. Ozer, **P. Lotfi-Kamran**, A. Panteli, A. Prodromou, C. Nicopoulos, D. Hardy, B. Falsafi, and Y. Sazeides, "Thermal Characterization of Cloud Workloads on a Power-Efficient Server-on-Chip," in *International Conference on Computer Design (ICCD)*, pp. 175–182, October 2012.
10. **P. Lotfi-Kamran**, B. Grot, M. Ferdman, S. Volos, O. Kocberber, J. Picorel, A. Adileh, D. Jevdjic, S. Idgunji, E. Ozer, and B. Falsafi, "Scale-Out Processors," in *International Symposium on Computer Architecture (ISCA)*, pp. 500–511, June 2012.
11. M. Hosseinabady, **P. Lotfi-Kamran**, J. Mathew, S. Mohanty, and D. Pradhan, "Single-Event Transient Analysis in High Speed Circuits," in *International Symposium on Electronic System Design (ISED)*, pp. 112–117, December 2011.
12. M. Ferdman, **P. Lotfi-Kamran**, K. Balet, and B. Falsafi, "Cuckoo Directory: A Scalable Directory for Many-Core Systems," in *International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 169–180, February 2011. (**Selected by the program committee for the best student paper session**)

13. **P. Lotfi-Kamran**, M. Ferdman, D. Crisan, and B. Falsafi, "TurboTag: Lookup Filtering to Reduce Coherence Directory Power," in *International Symposium on Low Power Electronics and Design (ISLPED)*, pp. 377–382, August 2010.
14. A.-M. Rahmani, I. Kamali, **P. Lotfi-Kamran**, A. Afzali-Kusha, and S. Safari, "Negative Exponential Distribution Traffic Pattern for Power/Performance Analysis of Network on Chips," in *International Conference on VLSI Design (VLSID)*, pp. 157–162, January 2009.
15. **P. Lotfi-Kamran**, M. Daneshtalab, C. Lucas, and Z. Navabi, "BARP–A Dynamic Routing Protocol for Balanced Distribution of Traffic in NOCs," in *Design, Automation and Test in Europe (DATE)*, pp. 541–546, March 2008.
16. **P. Lotfi-Kamran**, A.-A. Salehpour, A.-M. Rahmani, A. Afzali-Kusha, and Z. Navabi "Stall Power Reduction in Pipelined Architecture Processors," in *International Conference on VLSI Design (VLSID)*, pp. 541–546, January 2008.
17. **P. Lotfi-Kamran**, M. Massoumi, M. Mirzaei, and Z. Navabi, "Enhanced TED: A New Data Structure for RTL Verification," in *International Conference on VLSI Design (VLSID)*, pp. 481–486, January 2008.
18. M. Hosseinabady, M. H. Neishaburi, **P. Lotfi-Kamran**, and Z. Navabi, "A UML Based System Level Failure Rate Assessment Technique for SoC Designs," in *VLSI Test Symposium (VTS)*, pp. 243–247, May 2007.
19. M. Hosseinabady, **P. Lotfi-Kamran**, G. Di Natale, S. Di Carlo, A. Benso, and P. Prinetto, "Single-Event Upset Analysis and Protection in High Speed Circuits," in *European Test Symposium (ETS)*, pp. 29–34, May 2006.
20. M. Hosseinabady, **P. Lotfi-Kamran**, P. Riahi, F. Lombardi, and Z. Navabi, "A Flow Graph Technique for DFT Controller Modification," in *International System-on-Chip Conference (SOCC)*, pp. 55–60, September 2005.
21. A. Hooshmand, S. Shamshiri, M. Alisafaee, B. Alizadeh, **P. Lotfi-Kamran**, M. Naderi, and Z. Navabi, "Binary Taylor Diagrams: An Efficient Implementation of Taylor Expansion Diagrams," in *International Symposium on Circuits and Systems (ISCAS)*, vol. 1, pp. 424–427, May 2005.
22. **P. Lotfi-Kamran**, H. Shojaei, H. Parandeh-Afshar, M. Naderi, and Z. Navabi, "Improving Logic-Level Representation of BMD/TED Diagrams," in *Iranian Conference on Electrical Engineering (ICEE)*, pp. 448–453, May 2005. (**Recognized as best student paper by the program committee**)
23. **P. Lotfi-Kamran**, M. Hosseinabady, H. Shojaei, M. Massoumi, and Z. Navabi, "TED+: A Data Structure for Microprocessor Verification," in *Asia South Pacific Design Automation Conference (ASP–DAC)*, vol. 1, pp. 567–572, January 2005.

Journal Papers

1. F. Golshan, M. Bakhshalipour, M. Shakerinava, A. Ansari, **P. Lotfi-Kamran**, and H. Sarbazi-Azad, "Harnessing Pairwise-Correlating Data Prefetching with Runahead Metadata," in *IEEE Computer Architecture Letters (CAL)*, vol. 19, no. 2, pp. 130–133, July-December 2020.
2. A. Ansari, **P. Lotfi-Kamran**, and H. Sarbazi-Azad, "Code Layout Optimization for Near-Ideal Instruction Cache," in *IEEE Computer Architecture Letters (CAL)*, vol. 18, no. 2, pp. 124–127, July-December 2019.
3. M. Bakhshalipour, S. Tabaeiaghdaei, **P. Lotfi-Kamran**, and H. Sarbazi-Azad, "Evaluation of Hardware Data Prefetchers on Server Processors," in *ACM Computing Surveys (CSUR)*, vol. 52, no. 3, article 52, June 2019.
4. M. Bakhshalipour, A. Faraji, S. A. Vakil Ghahani, F. Samandi, **P. Lotfi-Kamran**, and H. Sarbazi-Azad, "Reducing Writebacks Through In-Cache Displacement," in *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, vol. 24, no. 2, article 16, March 2019.
5. M. Bakhshalipour, **P. Lotfi-Kamran**, A. Mazloumi, F. Samandi, M. Naderan-Tahan, M. Modarresi, and H. Sarbazi-Azad, "Fast Data Delivery for Many-Core Processors," in *IEEE Transactions on Computers (TC)*, vol. 67, no. 10, pp. 1416–1429, October 2018.
6. A. Vakil-Ghahani, S. Mahdizadeh-Shahri, M. Lotfi-Namin, M. Bakhshalipour, **P. Lotfi-Kamran**, and H. Sarbazi-Azad, "Cache Replacement Policy Based on Expected Hit Count," in *IEEE Computer Architecture Letters (CAL)*, vol. 17, no. 1, pp. 64–67, January-June 2018.
7. M. Bakhshalipour, **P. Lotfi-Kamran**, and H. Sarbazi-Azad, "An Efficient Temporal Data Prefetcher for L1 Caches," in *IEEE Computer Architecture Letters (CAL)*, vol. 16, no. 2, pp. 99–102, July-December 2017.

8. H. Sarbazi-Azad and **P. Lotfi-Kamran**, "Inefficiencies of Evaluation Mechanism for Promotion and Tenure of Researchers and Faculty Members in Iran," in *The CSI Journal on Computing Science and Information Technology (CSIT)*, vol. 15, no. 1, pp. 1–9, October 2017.
9. A. Yazdanbakhsh, D. Mahajan, H. Esmaeilzadeh, and **P. Lotfi-Kamran**, "AxBench: A Multi-Platform Benchmark Suite for Approximate Computing," in *IEEE Design & Test (D&T), Special Issue on Computing in the Dark Silicon Era*, vol. 34, no. 2, pp. 60–68, April 2017.
10. S. H. Seyyedaghaei Rezaei, A. Mazloumi, M. Modarressi, and **P. Lotfi-Kamran**, "Dynamic Resource Sharing for High-Performance 3-D Networks-on-Chip," in *IEEE Computer Architecture Letters (CAL)*, vol. 15, no. 1, pp. 5–8, January–June 2016.
11. **P. Lotfi-Kamran**, M. Modarressi, and H. Sarbazi-Azad, "An Efficient Hybrid-Switched Network-on-Chip for Chip Multiprocessors," in *IEEE Transactions on Computers (TC)*, vol. 65, no. 5, pp. 1656–1662, May 2016.
12. **P. Lotfi-Kamran**, "Per-Packet Global Congestion Estimation for Fast Packet Delivery in Networks-on-Chip," in *The Journal of Supercomputing (SUPE)*, vol. 71, no. 9, pp. 3419–3439, September 2015.
13. B. Grot, D. Hardy, **P. Lotfi-Kamran**, B. Falsafi, C. Nicopoulos, and Y. Sazeides, "Optimizing Data-Center TCO with Scale-Out Processors," in *IEEE Micro, Special Issue on Energy-Aware Computing*, vol. 32, no. 5, pp. 52–63, September–October 2012.
14. **P. Lotfi-Kamran**, A.-M. Rahmani, M. Daneshtalab, A. Afzali-Kusha, and Z. Navabi, "EDXY–A Low-Cost Congestion-Aware Routing Algorithm for Network-on-Chips," in *Elsevier Journal of Systems Architecture – Embedded Systems Design (JSA–ESD)*, vol. 56, no. 7, pp. 256–264, July 2010.
15. M. Hosseinabady, **P. Lotfi-Kamran**, F. Lombardi, and Z. Navabi, "Low Overhead DFT Using CDFG by Modifying Controller," in *IET Computers & Digital Techniques (IET–CDT)*, vol. 1, no. 4, pp. 322–333, July 2007.
16. M. Hosseinabady, **P. Lotfi-Kamran**, and Z. Navabi, "Low Test Application Time Resource Binding for Behavioral Synthesis," in *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, vol. 12, no. 2, article 16, April 2007.
17. **P. Lotfi-Kamran** and Z. Navabi, "Improving Logic-Level Representation of Taylor Expansion Diagram Using Attributed Edges," in *Iranian Journal of Science and Technology (IJST), Transaction B, Engineering*, vol. 30, no. B6, pp. 735–748, December 2006.

Book Chapters

1. **P. Lotfi-Kamran** and H. Sarbazi-Azad, "Dark Silicon and the History of Computing," *Chapter 1, Advances in Computers: Dark Silicon and Future On-chip Systems, Elsevier Inc.*, vol. 110, pp. 1–33, July 2018.
2. A. Yazdanbakhsh, J. Park, H. Sharma, **P. Lotfi-Kamran**, and H. Esmaeilzadeh, "Accelerating GPU Accelerators Through Neural Algorithmic Transformation," *Chapter 16, Advances in GPU Research and Practice, Elsevier Inc.*, pp. 417–448, September 2016.

Workshop & Poster Papers

1. M. Dashtbani, **P. Lotfi-Kamran**, D. Rahmati, S. Gorgin, and H. Sarbazi-Azad, "Iran's National Grid Initiative: Objectives, Challenges, and Opportunities," in *Big Data and HPC: Ecosystem and Convergence, Russian Supercomputing Day (RUSCDAYS)*, September 2017.
2. A. Yazdanbakhsh, J. Sacks, C. Song, **P. Lotfi-Kamran**, H. Esmaeilzadeh, and N. Sung Kim, "NAX: Near-Data Approximate Computing," in *Workshop on Approximate Computing (AC)*, October 2016.
3. S. Idgunji, D. Milojevic, E. Ozer, **P. Lotfi-Kamran**, D. Jevdjic, and B. Falsafi, "Performance and Efficiency of 3D-Stacked DRAM in a Multicore System," in *Workshop on 3D Integration – Applications, Technology, Architecture, Design, Automation and Test in Conjunction with Design, Automation and Test in Europe (DATE)*, March 2012.
4. E. Ozer, K. Flautner, S. Idgunji, A. Saidi, Y. Sazeides, B. Ahsan, N. Ladas, C. Nicopoulos, I. Sideris, B. Falsafi, A. Adileh, M. Ferdman, **P. Lotfi-Kamran**, M. Kuulusa, P. Marchal, and N. Minas, "EuroCloud: Energy-Conscious 3D Server-on-Chip for Green Cloud Services," in *Workshop on Architectural Concerns in Large Datacenters in Conjunction with International Symposium on Computer Architecture (ISCA)*, June 2010.

5. M. Hosseinabady, **P. Lotfi-Kamran**, and Z. Navabi, "Controller-Aware Hierarchical Test Generation and Testability Analysis," in *European Test Symposium (ETS)*, May 2005.
6. M. Alisafaee, **P. Lotfi-Kamran**, S. Shamshiri, H. Esmaeilzadeh, A. Pedram, and Z. Navabi, "MCBIST: A New Online BIST Scheme," in *Workshop on RTL and High Level Testing (WRTLT)*, pp. 85–90, November 2004.
7. M. Hosseinabady, **P. Lotfi-Kamran**, A. Pedram, and Z. Navabi, "A Binary Wavelet Test Compression," in *Workshop on RTL and High Level Testing (WRTLT)*, November 2004.
8. S. Shamshiri, H. Esmaeilzadeh, M. Alisafaee, **P. Lotfi-Kamran**, and Z. Navabi, "Test Instruction Set (TIS): An Instruction Level CPU Core Self-Testing Method," in *European Test Symposium (ETS)*, May 2004.

Patent

1. B. Falsafi, B. Grot, and **P. Lotfi-Kamran**, "Network-on-Chip using Request and Reply Trees for Low-Latency Processor-Memory Communication," US Patent 9,703,707, Filed December 4, 2012, Issued July 11, 2017.

Research Experience

- **Principal investigator.** *School of Computer Science* Jul. 2014–date
  Institute for Research in Fundamental Sciences (IPM)

  Divide and Conquer Frontend Bottleneck. *The frontend stalls caused by instruction and BTB misses are a significant source of performance degradation in server processors. Prefetchers are commonly employed to mitigate frontend bottleneck. However, next-line prefetchers, which are available in server processors, are incapable of eliminating a considerable number of L1 instruction misses. Temporal instruction prefetchers, on the other hand, effectively remove most of the instruction and BTB misses but impose significant area overhead. Recently, an old idea of using BTB-directed instruction prefetching is revived to address the limitations of temporal instruction prefetchers. While this approach leads to prefetchers with low area overhead, it requires significant changes to the frontend of a processor. Moreover, as this approach relies on the BTB content for prefetching, BTB misses stall the prefetcher, and likely lead to costly instruction misses. In this work, we divide the frontend bottleneck into three categories and use a divide-and-conquer approach to propose simple and effective solutions for each one. Our proposal, SN4L+Dis+BTB, imposes the same area overhead as the state-of-the-art BTB-directed prefetcher, and at the same time, outperforms it by 5% on average and up to 16% [ISCA'20].*

  Bingo Spatial Data Prefetcher. *Applications extensively use data objects with a regular and fixed layout, which leads to recurrence of access patterns over memory regions. Spatial data prefetchers exploit this phenomenon to prefetch future memory references. They associate observed access patterns to an event. Upon recurrence of the same event, the recorded access pattern is used for prefetching. Existing spatial prefetchers associate observed access patterns to either a long sequence of events (e.g., A, then B, then C, ...) with low probability of recurrence or a short sequence of events (e.g., A) with high probability of recurrence. As a result, the prefetchers either offer low accuracy or lose significant prediction opportunities. We make a case for Bingo prefetcher that associates the observed spatial patterns to both short and long events to achieve high accuracy and not lose prediction opportunities. Through a detailed evaluation, we show that Bingo improves system performance by 16% over the state of the art [HPCA'19].*

  Domino Temporal Data Prefetcher. *Big-data applications frequently encounter data misses, and hence, lose significant performance potential. While state-of-the-art temporal prefetching techniques are effective at reducing the number of data misses, we observe that there is a significant gap between what they offer and the opportunity. Temporal prefetchers lookup the recorded history of past cache misses to choose a sequence of misses for prefetching. We find that the lookup mechanism of existing temporal prefetchers is responsible for the large gap. Existing lookup mechanisms either not choose the right sequence of misses in the history, or unnecessarily delay the sequence selection, and hence, miss the opportunity at the beginning of every sequence. To address such limitations, Domino prefetcher logically looks up the history with both one and two last miss addresses to find a match for prefetching. We propose a practical design for Domino prefetcher and show that it captures more than 90% of the temporal opportunity [CAL'17, HPCA'18].*

  Near-Ideal 2D NOCs for Servers. *Most of the NOC delay is due to per-hop resource allocation. In this work, we take advantage of proactive resource allocation (PRA) to eliminate per-hop resource allocation time in single-cycle multi-hop networks to reach a near-ideal network for servers. PRA is undertaken during (1) the time interval in which it is known that LLC has the requested data, but the data is not yet ready, and (2) the time interval in which a packet is stalled in a router because the required resources are dedicated to another packet. Through detailed evaluation, we showed that PRA improves system performance by 12% over the state-of-the-art single-cycle multi-hop mesh NOC, and is behind an ideal network by less than 4% [HPCA'17].*

Resource Sharing for 3D NOCs. *To turn the 3D technology bandwidth and latency benefits into network latency reductions and performance improvement, we need networks-on-chip (NOCs) that are specially designed to take advantage of what 3D technology has to offer. While in parallel workloads many packets experience blocking in the network due to losing arbitration for crossbars' input/output ports, we observe that in a considerable fraction of these cases in a 3D NOC, the corresponding input and output ports of the crossbar in the above or below router are idle. Given this observation, we proposed FRESH, a router microarchitecture with Fine-grained 3D REsource SHaring capability that leverages the ultra-low latency vertical links of a 3D chip to share crossbars and links at a fine granularity between vertically stacked routers. It enables packets that lose arbitration for crossbars' input/output ports to use idle resources of the above or below routers, and effectively eliminates the unnecessary packet blocking time. We showed that our proposal lowers network latency by up to 21% over the state-of-the-art 3D NOC [CAL'16].*

Neural Accelerators for Graphics Processing Units. *Graphics processing units (GPUs) accelerate the execution of diverse classes of applications, such as recognition, gaming, data analytics, weather prediction, and multimedia. Many of these applications are amenable to approximate execution. This application characteristic provides an opportunity to improve the performance and efficiency of GPUs. We studied the effectiveness of neural approximate acceleration for GPU workloads. We showed that applying CPU neural accelerators to GPUs leads to high area overhead. Therefore, we defined a low overhead neurally accelerated architecture for GPUs that enables scalable integration of neural acceleration on a large number of GPU cores. Our evaluation indicated 2.4× speedup and 2.8× energy reduction with 10% quality loss and 1.2% area overhead [MICRO'15].*

Benchmark Suite for Approximate Computing. *Approximate computing promises to deliver significant performance and energy efficiency gains when small losses of quality are permissible. As approximate computing attracts more attention, having a general, diverse, and representative set of benchmarks to evaluate different approximation techniques becomes necessary. In this work, we introduced AxBench, a general, diverse and representative set of benchmarks for CPUs, GPUs, and hardware design [D&T'17].*

- **Research assistant.** *Parallel Systems Architecture (PARSA) Lab* Sep. 2008–Sep. 2013
  Swiss Federal Institute of Technology in Lausanne (EPFL)

  Advisor: Babak Falsafi

  Organization of Scale-Out Processors. *I proposed NOC-Out, a many-core organization for Scale-Out Processors that has low area overhead and provides fast access to the last-level cache (LLC) for delivering high performance. Today's many-core organizations force a compromise between performance and cost. Thus, many-core organizations based on a mesh interconnect have a modest area and wire cost, yet incur latency overheads through a many-hop topology. In contrast, many-core organizations based on richly connected topologies, such as a flattened butterfly, offer low latency at high area and wire cost. While existing many-core organizations offer an uneasy compromise between low area overhead and fast access to the LLC, NOC-Out offers both features simultaneously. The proposed organization is based on one simple and critical observation: there is almost no core-to-core communication in scale-out workloads. Based on this observation, this organization decouples cores and the last-level cache, eliminates all unneeded core-to-core links, and uses specialized core-to-LLC networks to connect cores to the last-level cache and vice versa. The bottom line is that NOC-Out delivers the performance of the state-of-the-art many-core organization with 1/10th of the area [MICRO'12].*

  Scale-Out Processors. *I proposed a methodology for the design of highly efficient many-core processors for scale-out workloads. This research relies on two critical observations with regard to such many-core processors. First, large LLCs waste precious silicon real estate that could have been better used to integrate more cores. Second, the organization of a many-core processor has a significant impact on its performance. Existing many-core chips, such as those offered by Tilera, sacrifice much of the on-die real estate to LLC and employ a tiled organization that incurs a high on-chip communication overhead. In contrast, I proposed a many-core processor based on the notion of pods. A pod is a module that tightly couples many cores to a modestly sized LLC through a low-latency interconnect. The proposed processor integrates many pods wherein each pod is a self-contained server-on-a-chip running a full software stack. I formulated a methodology to determine the optimal number of cores and LLC capacity to integrate in a pod. The proposed design, called the Scale-Out Processor, delivers peak throughput in today's process technology and affords near-ideal scalability as the technology scales [ISCA'12].*

Area- and Energy-Efficient Coherence Directories for Many-Core CMPs. *We proposed the Cuckoo directory, an energy- and area-efficient scalable directory organization. Existing directory organizations suffer from the lack of energy or area efficiency at high core counts due to wide associative lookups or capacity overprovisioning. The Cuckoo directory, however, scales to high core counts without the energy costs of wide associative lookup and without gross capacity overprovisioning. The Cuckoo directory borrows heavily from Cuckoo Hashing, a dense-storage software hashing technique. Rather than significantly over-provisioning storage capacity to avoid storage conflicts in a traditional lookup table, the Cuckoo directory uses the Cuckoo Hashing algorithm to relocate conflicting entries within the directory to alternate non-conflicting locations. Leveraging the mathematically robust properties of Cuckoo Hashing enables compact and energy-efficient coherence directories with predictable asymptotic behavior and without degenerate cases, improving the state-of-the-art directory design without increasing complexity [HPCA'11].*

Lookup Filtering to Reduce Power Consumption of Coherence Directories. *To reduce the energy usage of coherence directories in many-core processors, I proposed the TurboTag filtering mechanism. Coherence directories dissipate a significant fraction of their power on unnecessary lookups when running commercial server and scientific workloads. These workloads have large working sets that are beyond the reach of on-chip caches of modern processors. Limited to capturing a small part of the working set, private caches retain cache blocks only for a short period of time before replacing them with new blocks. Moreover, coherence enforcement is a known performance bottleneck of multi-threaded software; hence, data sharing in optimized high-performance software is minimal. Consequently, the majority of the accesses to the coherence directory find no sharers in the directory because the data are not available in the on-chip private caches, effectively wasting power on the coherence checks. TurboTag reduces power consumption of coherence directories by filtering (almost all) needless directory lookups [ISLPED'10].*

CloudSuite Workloads. *CloudSuite is a benchmark suite for emerging scale-out applications. To enable full-system simulation of CloudSuite benchmarks, we brought up CloudSuite workloads on Flexus, our in-house, full-system simulator. I led the effort to bring up CloudSuite workloads. We released the workloads to the broader research community.*

Flexus Full-System Simulator. *Flexus is a family of component-based C++ computer architecture simulators that enable full-system, timing-accurate simulation of uni- and multi-processor systems running unmodified commercial applications and operating systems. I contributed to the development of Flexus and also served as the coordinator of its mailing list.*

- **Research assistant.** *Computer-Aided Design (CAD) Lab* Jul. 2002–Aug. 2008
  University of Tehran

  Advisor: Zainalabedin Navabi

  Cost-Effective Globally Aware Dynamic Routing Protocols to Avoid Congestion in NOCs. *I proposed low-cost adaptive routing algorithms for network-wide congestion avoidance. This research enhanced conventional adaptive routings that only rely on local indicators for congestion estimation with low-cost, non-local (global) indicators [DATE'08, JSA–ESD'10].*

  Statistical Analysis for Soft Error Rate Estimation. *The effect of Single-Event Transients (SETs) on the system reliability is a big concern for ICs manufactured using advanced technologies. An SET in the combinational part of a circuit may propagate as a transient pulse to the input of a flip-flop and, consequently, gets latched, thus generating a soft error. Using the Probability Density Function (PDF) of an SET, we proposed a statistical method to compute the probability of soft errors considering dynamic behavior of a circuit [ETS'06, ISED'11].*

  Low-Overhead Controller-Aware Design for Test. *We proposed a low-cost design for test (DFT) that requires minor modifications to the controller of digital systems. In this research, we took advantage of existing data paths in digital systems to provide controllability and observability for the test process. Furthermore, we introduced additional data paths by altering states or adding new transitions in the controllers of digital systems. The proposed DFT considerably reduces the test application time by ignoring unnecessary control states in the test process [IET–CDT'07].*

  Low Test Application Time Test Synthesis. *Increased density and the need to test for new types of defects in nanometer technologies have resulted in a tremendous increase in test application time. We presented a test synthesis mechanism to reduce test application time for testing the datapath of a digital system. The reduction in the test application time was achieved by applying a test time-aware, resource-sharing algorithm on a scheduled control data flow graph of a design [TODAES'07].*

  Data Structure for Efficient RT-Level Representation. *Formal verification of microprocessors requires a mechanism for efficient representation and manipulation of both arithmetic and random Boolean functions. State-of-the-art representations can effectively represent arithmetic expressions at the word-level but are not memory efficient in representing bit-level logic expressions. In this research, I presented modifications to the state-of-the-art representations to improve the ability of bit-level logic representation while maintaining robustness in arithmetic word-level representation [ASP–DAC'05].*

VHDL-AMS Analyzer. *VHDL-AMS is a derivative of the hardware description language VHDL that includes analog and mixed-signal extensions (AMS) in order to define the behavior of analog and mixed-signal systems. I developed an analyzer to compile VHDL-AMS to CHIRE, our in-house intermediate format.*

Teaching Experience

- **Instructor.** Designed course, gave lectures, held office hours, answered email/newsgroup queries, designed and graded homework and exams.
  - CE237: Advanced Multiprocessor Architecture. Iran University of Science & Technology — Fall 2018, Spring 2020, Spring 2021
  - CE723: Advanced Computer Architecture. Iran University of Science & Technology — Fall 2020
  - CE723: Advanced Computer Architecture. Sharif University of Technology — Spring 2014–Fall 2014, Spring 2016–Fall 2016
  - CE323: Introduction to Computer Architecture. Sharif University of Technology — Spring 2015, Spring 2016
  - CE722: Advanced Microarchitecture. Sharif University of Technology — Fall 2015
  - ECE061: Discrete Mathematical Structures. University of Tehran — Spring 2007–Spring 2008
  - Operating Systems. University of Tehran, IT education series for professionals with no computer science background. — Spring 2003
- **Teaching assistant.** Held office hours, answered email/newsgroup queries, led review sessions, designed and graded student homework.
  - CS471: Advanced Multiprocessor Architecture. EPFL — Fall 2009 and Fall 2012
  - CS198: Information Technology Project. EPFL — Fall 2011
  - ECE367: Digital Logic Circuits. University of Tehran — Spring 2004–Spring 2005
  - ECE354: Telecommunications 1. University of Tehran — Fall 2003–Spring 2004
  - ECE046: Microprocessors Lab. University of Tehran — Spring 2003
  - ECE207: Microprocessors. University of Tehran — Spring 2002

Professional Activities

- **PC and ERC member.**
  - International Symposium on Microarchitecture (MICRO) — 2021
  - International Symposium on Computer Architecture (ISCA) — 2021
  - International Symposium on Computer Architecture (ISCA) — 2020
- **Reviewer.**
  - ACM Transactions on Computer Systems (TOCS) — 2016-2021
  - IEEE Computer Architecture Letters (CAL) — 2015-2021
  - IEEE Transactions on Computers (TC) — 2014-2021
  - International Symposium on High-Performance Computer Architecture (HPCA) — 2013-2014
  - ACM Transactions on Architecture and Code Optimization (TACO) — 2012-2021
  - Elsevier Journal of Computer and System Sciences (JCSS) — 2014
  - ACM Transactions on Embedded Computing Systems (TECS) — 2013
  - International Symposium on Microarchitecture (MICRO) — 2012
  - Elsevier Journal of System Architecture (JSA) — 2012
- **Member.**
  - ACM and ACM SIGARCH — Nov. 2008–date
  - IEEE and IEEE Computer Society — Jan. 2005–date

Tutorials

- **CloudSuite on Flexus.**
  - CSICC 2013, Tehran, Tehran, Iran — Mar. 2013
  - ISCA 2012, Portland, Oregon, USA — Jun. 2012
- **Flexus.**
  - IISWC 2010, Atlanta, Georgia, USA — Dec. 2010

Work Experience

- **Senior Design Engineer.** Public Switched Telephone Network (PSTN) Branch, Parstel Information and Telecommunication Technology Co. Inc. — Dec. 2005–Aug. 2008
- **Design Engineer.** Programmable Logic Controller (PLC) Branch, Arman Optimized Systems — Jun. 2001–Feb. 2002