

AMIN FIROOZSHAHIAN

Phone: (650) 714-0592

aminf13@gmail.com

www.firoozshahian.com

Mountain View, CA 94040

EXPERIENCE

Rain AI, San Francisco, CA

May 2024 – Present

Lead Architect

- Leading architecture team
- Driving architecture development of the next generation energy efficient AI accelerators

Meta Platforms Inc., Menlo Park, CA

Feb. 2019 – May 2024

Research Scientist, Infrastructure

- Leading architect, core ASIC architecture team
 - Leading a team of experts, defining, specifying and documenting next generation deep learning accelerator architectures
- Leading architectural model development
 - Functional models used for software prototyping
 - Cycle-accurate models used for performance evaluation and design verification
- Leading vendor interaction for processor (RISC-V) IP procurement
 - Specifying IP requirements and expectations
 - Configuring, customizing and deploying IPs internally
 - Handling internal maintenance and support

Intel Corporation, Santa Clara, CA

Nov. 2014 – Feb. 2019

Senior Staff - Silicon Architecture Engineer, Product Architecture Group

- Led architecture development and software enablement of novel technology for Xeon family of products
 - Drove definition of the architecture and hardware/software interfaces
 - Collaborated with various software teams for technology enablement
 - Closely cooperated with the micro-architecture and design teams for technology implementation and validation
 - Lead performance evaluation and benchmarking of various aspects of the technology
- Member of the AIPG performance evaluation team
 - Developed cycle-accurate performance simulator of the Crest product line for deep learning acceleration
- Member of the Architecture Patent Committee
 - Evaluated and voted on architecture related patent proposals

Hicamp Systems Inc, Menlo Park, CA

Oct. 2008 – Nov. 2014

Member of Technical Staff

- One of the first two employees joining the company
- Key member of the technical team developing and prototyping novel technology

- Owned and delivered two major blocks of the architecture
- Responsibilities included:
 - Functional simulation
 - Architecture definition
 - Micro-architecture design
 - RTL coding and implementation
 - Test and verification
 - FPGA synthesis and timing closure
 - System debug and bring up
- Performed benchmarking and performance evaluation of the technology on the prototype system
- ***Acquired by Intel in 2014***

Stanford University, Stanford, CA

Jan. 2001 – Oct. 2008

Research Assistant, Computer Systems Laboratory

- Architect for Smart Memories Project
- Designed a reconfigurable memory system for the test chip, with support for cache coherence, streaming and transactional memory models
- Developed a functional simulator to evaluate performance of the memory system architecture
- Designed micro-architecture and developed RTL implementation of a reconfigurable protocol controller for supporting cache coherent, streaming and transactional memory models
- Synthesized, and performed pre-silicon test and verification of Smart Memories multiprocessor system
- Designed, implemented and tested an 8×8 packet switch with support for eight virtual channels and broadcasting for connecting Smart Memories test chips on a multi-chip board

Talesh Electronic Company (TEC), Tehran, Iran

1996 – 2001

Hardware/System Engineer

- Involved in development of a number of projects, including industrial automation, data acquisition and microprocessor-based systems

EDUCATION

PhD	Stanford University, Electrical Engineering Department Stanford, CA, 94305 Dissertation: “Smart Memories: A Reconfigurable Memory System Architecture”	Jan. 2009
MS	University of Tehran, Faculty of Engineering Tehran, Iran	Jan. 2001
BS	Sharif University of Technology, Computer Engineering Department Tehran, Iran	Jun. 1998

AWARDS AND HONORS

Winner of the Best Paper Award 26th International Conference on Supercomputing (ICS '12)	2012
--	------

SELECTED PUBLICATIONS

Journal Papers

M. Wachs, O. Sacham, Z. Asgar, A. Firoozshahian, S. Richardson and M. Horowitz, "[Bringing Up a Chip on the Cheap](#)," IEEE Design and Test of Computers, vol. 29, no. 6, Dec. 2012, pp. 57-65.

O. Sacham, O. Azizi, M. Wachs, W. Qadeer, A. Asgar, K. Kelley, J.P. Stevenson, S. Richardson, M. Horowitz, B. Lee, A. Solomatnikov and A. Firoozshahian, "[Rethinking Digital Design: Why Design Must Change](#)," IEEE Micro, volume 30, issue 6, Nov.-Dec. 2010, pp. 9-24 (*Invited Paper*).

J. Leverich, H. Arakida, A. Solomatnikov, A. Firoozshahian, M. Horowitz and C. Kozyrakis, "[Comparative Evaluation of Memory Models for Chip Multi-Processors](#)," ACM Transactions on Architecture and Code Optimization (TACO), vol. 5, no. 3, Nov. 2008, Article no. 12.

Conference Papers

Amin Firoozshahian, Joel Coburn, Roman Levenstein, Rakesh Nattoji, Ashwin Kamath, Olivia Wu, Gurdeepak Grewal, Harish Aepala, Bhasker Jakka, Bob Dreyer, Adam Hutchin, Utku Diril, Krishnakumar Nair, Ehsan K. Aredestani, Martin Schatz, Yuchen Hao, Rakesh Komuravelli, Kunming Ho, Sameer Abu Asal, Joe Shajrawi, Kevin Quinn, Nagesh Sreedhara, Pankaj Kansal, Willie Wei, Dheepak Jayaraman, Linda Cheng, Pritam Chopda, Eric Wang, Ajay Bikumandla, Arun Karthik Sengottuvel, Krishna Thottempudi, Ashwin Narasimha, Brian Dodds, Cao Gao, Jiyuan Zhang, Mohammed Al-Sanabani, Ana Zehtabioskuie, Jordan Fix, Hangchen Yu, Richard Li, Kaustubh Gondkar, Jack Montgomery, Mike Tsai, Saritha Dwarakapuram, Sanjay Desai, Nili Avidan, Poorvaja Ramani, Karthik Narayanan, Ajit Mathews, Sethu Gopal, Maxim Naumov, Vijay Rao, Krishna Noru, Harikrishna Reddy, Prahlad Venkatapuram, Alexis Bjorlin, "[MTIA: First Generation Silicon Targeting Meta's Recommendation Systems](#)," Proceedings of the 50th Annual International Symposium on Computer Architecture (ISCA '23), June 17-21, 2023, Orlando, FL.

L. Ke, U. Gupta, B. Y. Cho, D. Brooks, V. Chandra, U. Diril, A. Firoozshahian, K. Hazelwood, B. Jia, H.S. Lee, M. Li, B. Maher, D. Mudigere, M. Naumov, M. Schatz, M. Smelyanskiy, X. Wang, B. Reagen, C.J. Wu, M. Hempstead, X. Zhang, "[RecNMP: Accelerating Personalized Recommendation with Near-Memory Processing](#)," Proceedings of the 47th Annual International Symposium on Computer Architecture (ISCA '20), May 30th – June 3rd 2020, Valencia, Spain, pp. 790-803.

S. Ghorbani, Z. Yang, P.B. Godfrey, Y. Ganjali and A. Firoozshahian, "[DRILL: Micro Load Balancing for Low-Latency Data Center Networks](#)," Proceedings of the Conference of the ACM Special Interest Group on Data Communication (SIGCOMM '17), Aug. 21-25, 2017, Los Angeles, CA, pp. 225-238.

S. Ghorbani, B. Godfrey, Y. Ganjali and A. Firoozshahian, "[Micro Load Balancing in Data Centers with DRILL](#)," Proceedings of the 14th ACM Workshop on Hot Topics in Networks (HotNets XIV), Nov. 16-17, 2015, Philadelphia, PA, Article no. 17.

H. Litz, D. Cheriton, A. Firoozshahian, O. Azizi and J.P. Stevenson., "[SI-TM: Improving Transactional Memory Abort Rates through Snapshot Isolation](#)," Proceedings of the 19th international conference on Architectural support for programming languages and operating systems (ASPLOS 2014), Mar. 01-05, 2014, Salt Lake City, UT, pp. 383-398.

J.P. Stevenson, A. Firoozshahian, A. Solomatnikov, M. Horowitz and D. Cheriton, “[Sparse Matrix-Vector Multiply on HICAMP Architecture](#),” Proceedings of the 26th ACM international conference on Supercomputing (ICS '12), Jun. 25-29, 2012, Venice, Italy, pp. 195-204 (*Winner of the Best Paper Award*).

D. Cheriton, A. Firoozshahian, A. Solomatnikov, J.P. Stevenson and O. Azizi, “[HICAMP: Architectural Support for Efficient Concurrency-Safe Shared Structured Data Access](#),” Proceedings of the seventeenth international conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS XVII), Mar. 03-07, 2012, London, England, pp. 287-300.

A. Solomatnikov, A. Firoozshahian, O. Shacham, Z. Asgar, M. Wachs, W. Qadeer, S. Richardson and M. Horowitz, “[Using a Configurable Processor Generator for Computer Architecture Prototyping](#),” Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO 42), Dec. 12-16, 2009, New York, NY, pp. 358-369.

A. Firoozshahian, A. Solomatnikov, O. Shacham, Z. Asgar, S. Richardson, C. Kozyrakis and M. Horowitz, “[A Memory System Design Framework: Creating Smart Memories](#),” Proceedings of the 36th annual international symposium on Computer architecture (ISCA '09), June. 20-24, 2009, Austin, TX, pp. 406-417.

O. Shacham, M. Wachs, A. Solomatnikov, A. Firoozshahian, S. Richardson and M. Horowitz, “[Verification of Chip Multiprocessor Memory Systems Using A Relaxed Scoreboard](#),” Proceedings of the 41st annual IEEE/ACM International Symposium on Microarchitecture (MICRO 41), Nov. 08-12, 2008, Lake Como, Italy, pp. 294-305.

J. Leverich, H. Arakida, A. Solomatnikov, A. Firoozshahian, M. Horowitz and C. Kozyrakis, “[Comparing Memory Systems for Chip Multi-Processors](#),” Proceedings of the 34th annual international symposium on Computer architecture (ISCA '07), Jun. 09-13, 2007, San Diego, CA, pp. 358-368.

A. Firoozshahian, V. Manshadi, A. Goel and B. Prabhakar, “[Efficient, Fully Local Algorithms for CIOQ Switches](#),” Proceedings of the 26th IEEE International Conference on Computer Communications (INFOCOM 2007), May. 06-12, 2007, Barcelona, Spain.