CURRICULUM VITAE (CV)

Dr. Georgios Michelogiannakis

mihelog@lbl.gov

EDUCATION / EXPERIENCE

- Research Scientist (currently): Lawrence Berkeley National Laboratory. Computational research division. Started December 2014.
- Lecturer: University of California at Berkeley electrical engineering and computer science department. January to June 2016.
- Lecturer: Stanford University electrical engineering department. April to June 2013.
- Postdoctoral Fellow: Lawrence Berkeley National Laboratory. Computational research division. Ended October 2014.
- PhD: Electrical Engineering at Stanford University. Graduation date: August 2012.
- Intern: NVIDIA research. June to August 2011.
- MSc: Computer Science at the University of Crete. Graduation date: July 2007.
- BSc: Computer Science at the University of Crete. Graduation date: June 2005.

DISTINCTIONS

- Postdoctoral research fellowship for academic years 2012 – 2013 and 2013 – 2014.
- Stanford Graduate Fellowship: This is an honorary fellowship for three years awarded by Stanford University. This means I was placed in the top 5% of the electrical engineering applicants for 2007.
- University of Crete Drettakis distinction: achieved highest University of Crete Computer Science degree mark in year 2005. Graduated with the highest undergraduate degree grade. Best year grade in class, year 2002. Second best year grade in 2003 and 2004.
- BSc Degree studentship by the Greek State Studentship Foundation, years 2002 - 2004.

RESEARCH EXPERIENCE

- September 2012 to present: *Lawrence Berkeley National Laboratory.* In LBNL I am part of the computer architecture group and have participated in numerous proposals and projects. My research activities focus on various aspects of high performance computing, HPC networks, on-chip networks, task placement, prefetching, dataflow architectures, computer architecture, chip multiprocessors, memory hierarchy and scheduling, and the future of digital computing after the end of MOSFET scaling.
- September 2007 until August 2012: *Research assistant at Stanford University*. My research focused on many different techniques to design energy-efficient on-chip networks as well as improving allocation efficiency. I also developed a congestion control technique for supercomputer and datacenter networks. My supervisor was Prof. William J. Dally.
- Summer 2011: *Research intern at NVIDIA*. My manager was Dr. Steve Keckler. My internship focused on designing on-chip networks for future industrial chips. Also, I examined implementation decisions for the network and the rest of the system for

future chips. My work sparked interest and is the basis for further on-chip network research internally in NVIDIA.

- February 2004 until June 2007: *Research assistant at ICS-FORTH and the University of Crete*. My research focused on reducing on-chip network latency to make it comparable to that of dedicated wires. My primary advisor was Prof. Manolis G.H. Katevenis.

PATENTS

- **Michelogiannakis, G**.; Shalf, J.; "COLLECTIVE MEMORY TRANSFER DEVICES AND METHODS FOR MULTIPLE-CORE PROCESSORS". Pending.

PROPOSALS

- Donofrio, D.; Chan, C.; **Michelogiannakis, G**.; Shalf, J.; Unat, D.; Jung, M.; "CoDEx: A Hardware/Software Co-Design Environment for the Exascale Era". *Accepted.*
- **Michelogiannakis, G**.; Bokor, J.; Donofrio, D.; Shalf, J.; "Continuing Digital Computing Performance Scaling Post Moore's Law". *Accepted.*
- Shalf, J.; Donofrio, D.; **Michelogiannakis, G**.; Fatollahi-Fard, F.; "Test and Evaluation of EDA Tools for Superconducting logic". *Accepted*.

CONFERENCE PUBLICATIONS (H-index 11)

1. Unat, D.; Nguyen, T.; Zhang, W.; Nufail, M.; Basteem, B.; **Michelogiannakis, G**.; Almerg, A.; Shalf, J.; "TiDA: High-Level Programming Abstractions for Data Locality Management". ISC 2016: *International Supercomputing Conference*, June 2016.
2. Fatollahi-Fard, F.; Donofrio, D.; **Michelogiannakis, G**.; Shalf, J.; "OpenSoC Fabric: On-chip Network Generator", ISPASS 2016: *International Symposium on Performance Analysis of Systems and Software*, April 2016.
3. **Michelogiannakis, G**.; Shalf, J.; "Variable-Width Datapath for On-Chip Network Static Power Reduction", NOCS 2014: *International Symposium on Networks-on-Chip*, 17-19 September 2014.
4. **Michelogiannakis, G**.; Williams, S.; Williams, A.; Shalf, J.; "Collective Memory Transfers for Multi-Core Chips", ICS 2014: *International Conference on Supercomputing*, 10-13 June 2014.
5. **Michelogiannakis, G**.; Jiang, N.; Becker, D.; Dally, W.J.; "Channel Reservation Protocol for Over-Subscribed Channels and Destinations", *SC 2013*: *Conference on High Performance Computing Networking, Storage and Analysis*, 17-22 November 2013.
6. **Michelogiannakis, G**.; Li, X.S.; Bailey, D.H.; Shalf, J.; "Extending Summation Precision for Network Reduction Operations", *SBAC-PAD 2013*: *International Symposium on Computer Architecture and High Performance Computing*, 23-26 October, 2013.
7. Jiang, N.; Becker, D.; **Michelogiannakis, G**.; Balfour, J.; Towles, B.; Kim, J.; Dally, W.J.; "A Detailed and Flexible Cycle-Accurate Network-on-Chip Simulator". *ISPASS 2013*. *IEEE International Symposium on Performance Analysis of Systems and Software*, April 21-13, 2013.
8. Becker, D; Jiang, N; **Michelogiannakis, G**; Dally, W.J.; "Adaptive Backpressure: Efficient Buffer Management for On-Chip Networks", *ICCD 2012*. *30th IEEE International Conference on Computer Design*, September 30 – October 3, 2012.

9. Jiang, N.; Becker, D.; **Michelogiannakis, G**.; Dally, W.J.; "Network Congestion Avoidance through Speculative Reservation", *HPCA 2012*: *18th International Symposium on High Performance Computer Architecture*, 25-29 February, 2012.
10. **Michelogiannakis, G**.; Jiang, N.; Becker, D.; Dally, W.J.; "Packet Chaining: Efficient Single-Cycle Allocation for On-Chip Networks", *MICRO 2011*. *44th Annual IEEE/ACM International Symposium on Microarchitecture*, 3-7 December 2011.
11. **Michelogiannakis, G**.; Sanchez, D.; Dally, W.J.; Kozyrakis, C.; "Evaluating Bufferless Flow Control for On-chip Networks", *NOCS 2010: Fourth ACM/IEEE International Symposium on Networks-on-Chip*, pp. 9-16, 2-6 May 2010. **Best paper award candidate.**
12. **Michelogiannakis, G**.; Dally, W.J.; "Router Designs for Elastic Buffer On-Chip Networks", *SC 2009*: *Conference on Conference on High Performance Computing Networking, Storage and Analysis*, 14-20 November 2009.
13. **Michelogiannakis, G**.; Balfour, J.; Dally, W.J.; "Elastic-Buffer Flow Control for On-Chip Networks", *HPCA 2009: IEEE 15th International Symposium on High Performance Computer Architecture*, vol., no., pp.151-162, 14-18 February 2009.
14. Papaefstathiou, V.; Pnevmatikatos, D.; Marazakis, M.; Kalokairinos, G.; Ioannou, A.; Papamichael, M.; Kavadias, S.; **Michelogiannakis, G**.; Katevenis, M.; "Prototyping Efficient Interprocessor Communication Mechanisms", *SAMOS IC 2007*, July 16-19, 2007.
15. **Michelogiannakis, G**.; Pnevmatikatos, D.; Katevenis, M.; "Approaching Ideal NoC Latency with Pre-Configured Routes", *NOCS 2007: First ACM/IEEE International Symposium on Networks-on-Chip*, vol., no., pp.153-162, 7-9 May 2007.

JOURNAL PUBLICATIONS

16. **Michelogiannakis, G**.; Li, X.S.; Bailey, D.H.; Shalf, J.; "Extending Summation Precision for Network Reduction Operations", *Springer International Journal of Parallel Computing*, vol. 43 no. 6, pp. 1218-1243, 2015.
17. **Michelogiannakis, G**.; Dally, W.; "Elastic-buffer Flow Control for On-chip Networks", *IEEE Transactions on Computers*, vol. 62 no. 2, pp. 295-309, February 2013.
18. **Michelogiannakis, G**.; Jiang, N.; Becker, D.; Dally, W.; "Packet Chaining: Efficient Single-Cycle Allocation for On-Chip Networks", IEEE *Computer Architecture Letters*. June 2011.
19. **Michelogiannakis, G**.; Becker, D.U.; Dally, W.J.; "Evaluating Elastic Buffer and Wormhole Flow Control", *IEEE Transactions on Computers*, vol. 60, no. 6, pp. 896-903, June 2011.
20. Sanchez, D.; **Michelogiannakis, G**.; Kozyrakis, C.; "An Analysis of On-Chip Interconnection Networks for Large-Scale Chip Multiprocessors", ACM Transactions on Architecture and Code Optimization. May 2010.

BOOK CHAPTERS

21. **Michelogiannakis, G**.; "On-chip networks for modern large-scale chips". *Encyclopedia of Information Science and Technology*, 3rd edition. 2014. IGI global.
22. **Michelogiannakis, G**.; "State of the art and future trends of datacenter networks". *Encyclopedia of Information Science and Technology*, 3rd edition. 2014. IGI global.

WORKSHOP PUBLICATIONS

23. **Michelogiannakis, G.**; Donofrio, D.; Shalf, J.; "Modelling of Novel Transistors, Manufacturing Technologies, and Architectures to Preserve Digital Computing Performance Scaling". PMES 2016: *Post-Moore's Era Supercomputing (PMES) Workshop*, November, 2016.
24. **Michelogiannakis, G.**; Donofrio, D.; Shalf, J.; "Modeling of Novel Transistors, Manufacturing Technologies, and Architectures to Preserve Digital Computing Performance Scaling". MODSIM 2016: *Workshop on Modelling and Simulation of Systems and Applications*, August 2016.
25. Fatollahi-Fard, F.; Donofrio, D; **Michelogiannakis, G.**; Shalf, J.; "OpenSoC Fabric: On-Chip Network Generator". NoCArc 2014: *Workshop on network-on-chip architectures*, December 2014.
26. Fatollahi-Fard, F.; Donofrio, D; **Michelogiannakis, G.**; Shalf, J.; "OpenSoC Fabric: On-Chip Network Generator", MODSIM 2014: *Workshop on Modelling and Simulation of Systems and Applications*, 13 August, 2014.
27. Unat, D.; **Michelogiannakis, G.**; Shalf, J.; "The Role of Modelling in Locality Optimizations", MODSIM 2014: *Workshop on Modelling and Simulation of Systems and Applications*, 13 August, 2014.
28. **Michelogiannakis, G**.; "Hardware Support for Collective Memory Transfers", WOSC 2013: *Workshop on Optimizing Stencil Computations*, 27 October 2013.
29. Papaefstathiou, V.; Kalokairinos, G.; Ioannou, A.; Papamichael, M.; **Michelogiannakis, G**.; Kavadias, S.; Vlachos, E.; Pnevmatikatos, D.; Katevenios, M.G.H.; "An FPGA-based Prototyping Platform for Research in High-Speed Interprocessor Communication", *2nd HiPEAC Industrial Workshop*. 2007.
30. **Michelogiannakis, G**.; Katevenis, M.G.H.; "Approaching Ideal Latency with Pre-Configured (but Run-Time Re-Configurable) Routers", *1st HiPEAC Industrial Workshop*. 2006.

TECHNICAL REPORTS

31. **Michelogiannakis, G.**; Donofrio, D.; Shalf, J.; Bachan, J.; "Continuing the Scaling of Digital Computing Post Moore's Law", Lawrence Berkeley National Laboratory technical report. April 2016.
32. **Michelogiannakis, G**.; Shalf, J.; "Variable-Width Datapath for On-Chip Network Static Power Reduction", Lawrence Berkeley National Laboratory technical report. November 2013.
33. **Michelogiannakis, G**.; Williams, A.; Shalf, J.; "Collective Memory Transfers for Multi-Core Chips", Lawrence Berkeley national laboratory technical report. November 2013.
34. **Michelogiannakis, G**.; Jiang, N.; Becker, D.; Dally, W.J.; "Packet Chaining: Efficient Single-Cycle Allocation for On-Chip Networks", CVA memo 128. April 2011.
35. Jiang, N; **Michelogiannakis, G.**; Becker, D.; Towles, B.; Dally, W.D.; "Booksim 2.0 User Manual", March 2010.
36. **Michelogiannakis, G**.; Dally, W.J.; "Router Designs for Elastic Buffer On-Chip Networks", CVA memo 125. June 2009.
37. **Michelogiannakis, G**.; Balfour, J.; Dally, W.J.; "Elastic Buffer Networks-on-Chip", CVA memo 124. August 2008.
38. **Michelogiannakis, G**.; "Approaching Ideal NoC Latency with Pre-Configured Routes", ICS-FORTH TR-391-07-2007.
39. Papaefstathiou, V.; Kalokairinos, G.; Ioannou, A.; Papamichael, M.; **Michelogiannakis, G**.; Kavadias, S.; Vlachos, E.; Pnevmatikatos, D.; Katevenios,

M.G.H.; “Design and Implementation of a Multi-Gigabit NIC and a Scalable Buffered Crossbar Switch”, ICS-FORTH TR-376-04-2006.

THESES

40. PhD thesis: “Energy-Efficient Flow-Control for On-Chip Networks”. June 2012.
41. Master’s thesis: “Approaching Ideal NoC Latency with Pre-Configured Routes”. June 2007.
42. Undergraduate degree thesis: “IPIF to PCI Bridge Specification”. May 2005.

PAPERS IN PROGRESS / REVIEW

43. **Michelogiannakis, G**.; Jeremiah, W.; Kenny, J.; Ibrahim, K.; Shalf, J.; “Hierarchical task placement for task placement in HPC networks”.
44. **Michelogiannakis, G**.; Williams, S.; Shalf, J.; “LLCP: Last Level Collective Hardware Prefetching For Data-Parallel Applications”.
45. Mohammadiyaghi, P.; **Michelogiannakis, G.**; Gratz, P.; “Speculative Lock Allocation for Chip Multiprocessors”.

INVITED TALKS

- SC16 (panel and workshop), November 2016.
- MODSIM, August 2016.
- SC15 (panel), November 2015.
- NOCS (tutorial), September 2014.
- HPEC (BoF), September 2014.
- MODSIM, August 2014.
- ICS-FORTH, June 2014.
- ASPIRE retreat with UC Berkeley, May 2014.
- University of Southern California, April 2014.
- University of Tennessee, February 2014.
- SIAM Conference on Parallel Processing for Scientific Computing, February 2013 (panel member).
- MIT, December 2013.
- Workshop on Optimizing Stencil Computations, October 2013 (panel member).
- UC Santa Cruz, March 2013.
- Lawrence Berkeley National Laboratory, May 2012.
- SeaMicro technologies, April 2012.
- UC San Diego, March 2012.
- UC Irvine, February 2012.
- Arizona State University, January 2012.
- University of Connecticut, January 2012.
- ICS-FORTH, September 2008, September 2009 and December 2011.
- Pervasive parallel systems lab at Stanford, 2010 & 2011.
- National Technical University of Athens, CSLAB, September 2009.
- ST Microelectronics, October 2009.

CONFERENCE PAPER PRESENTATIONS

- 8th International Symposium on Networks-on-Chip (NOCS) 2014. Title: "Variable-Width Datapath for On-Chip Network Static Power Reduction".
- Workshop on Modelling and Simulation of Systems and Applications (MODSIM) 2014. Title: "The Role of Modelling in Locality Optimizations".
- 28th International ACM Conference on Supercomputing (ICS) 2014. Title: "Collective Memory Transfers for Multi-Core Chips".
- 25th Annual IEEE/ACM Conference on Conference on High Performance Computing Networking, Storage and Analysis (SC) 2013. Title: "Channel Reservation Protocol for Over-Subscribed Channels and Destinations".
- Workshop on Optimizing Stencil Computations (WOSC) 2013. Title: Hardware Support for Collective Memory Transfers in Stencil Computations". *Panel presentation*.
- 25th Annual IEEE International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD) 2013. Title: "Extending Summation Precision for Distributed Network Operations".
- 44th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO) 2011. Title: "Packet Chaining: Efficient Single-cycle Allocation for On-Chip Networks".
- 4th ACM/IEEE International Symposium on Networks-on-Chip (NOCS) 2010. Title: "Evaluating Bufferless Flow Control for On-Chip Networks".
- 21th Annual IEEE/ACM Conference on Conference on High Performance Computing Networking, Storage and Analysis (SC) 2009. Title: "Router Designs for Elastic Buffer On-Chip Networks".
- 15th Annual International Symposium on High Performance Computer Architecture (HPCA) 2009. Title: "Elastic Buffer Flow Control for On-Chip Networks".
- 1st ACM/IEEE International Symposium on Networks-on-Chip (NOCS) 2007. Title: "Approaching Ideal NoC Latency with Pre-Configured Routes".

TUTORIALS / PANELS

- SC 2016: Organizing a panel on the topic of technologies that will continue performance scaling of the classical digital computing model after MOSFETs stop scaling. The goal is to discuss a range of promising technologies and predict their state in a 20 year timeframe.
- SC 2015: Organized a panel with the title "Digital versus Neuromorphic versus Quantum". This panel brought together two experts from each of the three computing areas to discuss the shortcomings, potential, and current state of each computing platform. The panel had 150 or so attendees and had the room filled to capacity from start to finish.
- SC 2014: Organized a BoF session with the title "Evaluation infrastructure for future large-scale chips and HPC systems". This BoF focused on the challenges and fundamental tradeoffs in simulating and emulating future many-core chips and large-scale systems.
- NOCS 2014: Organized a tutorial with the title "OpenSoC: a Flexible, Parameterizable, Open NoC Generation Tool". This tutorial gave attendees a firm understanding as well as some hands-on experience on the infrastructure we are developing to generate both software and hardware models of on-chip networks and later entire SoCs in an open source manner.

CLASS TEACHING EXPERIENCE

- Spring 2016: Instructor of CS152 ("Computer architecture and Engineering") as a lecturer at UC Berkeley for the spring semester of 2016. There were 30 enrolled students. This is a senior-level undergraduate course. As the instructor, I was responsible for the lectures, compiling the exams and problem sets, partly grading, office hours, and mentoring the teaching assistant who helped with grading, and ran discussion sessions, labs, and office hours. This class resulted in three interns (two summer and one fall) for my group at LBNL. Students rated me the highest among the course's last few offerings, with particularly high scores in "gives lectures that are well-organized", "is enthusiastic about the subject matter", "encourages questions", "is helpful during office hours". My overall effectiveness was graded at 6.3/7 and the question "was this course worthwhile compared to others in UC" was graded 6.2/7.
- Spring 2013: Instructor of EE382c ("Interconnection networks") with the position of lecturer at Stanford for the spring quarter of 2013. There were 7 enrolled students. I was responsible for everything in the course, including preparing and giving lectures, making and grading homeworks and exams, mentoring and assisting students outside of lectures, as well as creating and supervising the final research project. The course was successful based on student's comments on their learning at the end of the course, and how it benefited students performing research on on-chip and system-wide networks. The scores of student evaluations reflected the successful outcome of this course. I was graded with an average score of 4.4/5.0 (20% higher than the department average) and a median score of 5.0/5.0. Particular areas of strength were "interaction with students", "ability to engage", and "class content".
- Spring 2011: Teaching assistant of EE382c at Stanford: "Interconnection networks". I supervised the research project, graded and wrote solutions for the homeworks and held office hours. I also prepared and gave six lectures of the class (one third of the total). The student evaluations reflected very well on my teaching, with an average score of 4.5/5 on various categories. The highest score was on teaching clarity.
- Spring 2007: Teaching assistant of CS255 at the Computer Science Department of the University of Crete: "Programming lab". I graded and supervised programming assignments and exams, and I held office hours. There were six assignments for the year, each stressing some aspect of object-oriented programming.
- Fall 2006 and Fall 2005: Teaching assistant of CS220 at the Computer Science Department of the University of Crete: "Digital circuits lab". I supervised lab sessions for four hours a week, during which I assisted and graded student's work. Students were tasked with implementing some logic in Verilog and testing it in FPGAs. I also graded exams.
- Spring 2006: Teaching assistant of CS325 at the Computer Science Department of the University of Crete: "Embedded systems". The other TA and I were tasked with designing and implementing the educational boards for the students. I supervised four to six hours of lab sessions each week, helped the students and graded them. I also graded exams.
- Ongoing: As an FAA and civil air patrol flight instructor, I regularly engage on close mentoring with student and established pilots. This type of mentoring is typically one-on-one. I have fully trained two private pilots and have provided various other training and evaluations to other pilots. I am currently working with two other students.

RESEARCH PROJECTS

- While at Lawrence Berkeley National Laboratory:
    - Continuing Digital Computing Performance Scaling After MOSFET Technology Stops Scaling: This project is critical for HPC and DOE given the large and important class of applications that depend on digital computing and do not fit the

quantum and neuromorphic computing models. The goal of this project is to first investigate emerging device, memory, and manufacturing technologies and quantify the impact they can have in continuing digital computing performance scaling. At the same time, this projects will look into specialized architectures for core HPC applications. At the same time, any drawbacks (e.g., reliability) need to be alleviated to enable quick adoption.

- Test and Evaluation of EDA Tools for Superconducting logic: The goal of this project is to test and evaluate computing logic that uses superconductors.

- Solid State Virtual Digital Fluid: This project designs a processor chip to structurally fit the problem (algorithm) it is designed for, and use manufacturing technologies expected in the next decade.

- Collective prefetcher: This project designs a hardware prefetcher that recognizes data access patterns created by data-parallel applications and predicts future data accesses that cores generate as a group. This prefetcher can span multiple memory pages without translations and requires a modest complexity increase compared to the strided prefetcher.

- Network Bandwidth Tapering Aware Task Placement: In future large-scale HPC and big data systems, bandwidth in each level of the network hierarchy will be limited by power, cost or even switch chip pin count. In fact, higher levels of the topology will most likely be more strictly constrained than lower levels, resulting in a topology with more bandwidth in the lower levels. In such a setting, it is critical that application threads or processes (such as MPI ranks) be placed such that communication avoids the higher levels of the topology as much as possible, in order to avoid bottlenecks but also minimize energy. This work utilizes graph bisection theory and applies it in a practical setting by investigating feasible algorithms that maximize communication locality. In addition, we apply this technique to large-scale systems without requiring user intervention, and examine the resulting reduced bandwidth requirements on the network.

- Speculative Lock Allocation for Chip Multiprocessors: This project attempts to exploit the sharing patterns of cache lines and locks in modern chip multiprocessors to assist data movement. It attacks performance bottlenecks created by overused cache lines and locks, which will become more critical with the simple cores projected for future many-core chips.

- TiDA: Tiling as a Durable Abstraction: Tiling is a useful loop transformation for expressing parallelism and data locality. Automated tiling transformations that preserve data-locality are increasingly important due to hardware trends towards massive parallelism and the increasing costs of data movement relative to the cost of computing. This work develops TiDA as a durable tiling abstraction that centralizes parameterized tiling information within array data types with minimal changes to the source code. The data layout information can be used by the compiler and runtime to automatically manage parallelism, optimize data locality, and schedule tasks intelligently.

- Collective memory transfers for multi-core chips: Performance improvements for a vast number of parallel applications depend on efficient access of neighboring data stored in main memory by a collection of cores on chip. A perfect example is the broad class of data parallel applications and kernels, which require domain-decomposition of data arrays from a contiguous arrangement in memory to a tiled layout for on-chip L1 data caches and scratchpads. However, DRAM

performance suffers under the non-streaming access patterns generated by many independent cores. DRAM performance and power are both crucial for a wide variety of chips. This project proposes collective memory scheduling (CMS) that actively takes control of collective memory transfers such that requests arrive in a sequential and predictable fashion to the memory controller. CMS maximizes memory throughput and reduces power beyond that attainable by even the most sophisticated memory controllers. This mechanism can be implemented in hardware-managed cache-coherent chip multiprocessors in the form of a last-level cache prefetcher [38], or with software assistance in architectures such as local stores using DMA operations [2].

- Extending summation precision for network reduction operations: Summation computational precision is at the core of numerous important algorithms such as Newton-Krylov methods and other operations involving inner products. To avoid precision loss and the accumulation of rounding errors in operations with millions or billions of operands as is typical in high-performance computing and big data processing, in this work I propose big integer (BigInt) expansions of double precision variables that enable arbitrarily large summations without error and provide exact and reproducible results for distributed (system-wide) summations. This is feasible with performance comparable to that of double-precision floating point summation and without the performance or cost penalties of sorting algorithms or complex data structures, by the inclusion of simple and inexpensive logic into modern NICs.

- Variable-width datapath for on-chip network static power reduction: With the tight power budgets in modern large-scale chips and the unpredictability of application traffic, on-chip network designers are faced with the dilemma of designing for worst-case bandwidth demands and incurring high static power overheads, or designing for an average traffic pattern and risk degrading performance. This work proposes adaptive bandwidth networks which reduce the static power of on-chip networks by dividing channels, buffers, and crossbars into lanes and activating only the number of lanes necessary in each hop to meet traffic demands. In addition, ABNs readily apply to silicon defect tolerance with just the extra cost for detecting faults.

- OpenSoC: a Flexible, Parameterizable, Open NoC Generation Tool: With future technology scaling, we anticipate many-core chips and large-scale systems to play a crucial role in applications ranging from HPC and big data down to embedded computing. However, the architectures we expect in the future are hard to simulate or emulate today. That is because detailed or cycle-accurate software simulation can be very slow, whereas hardware emulation requires an often excessively long development time and offers limited internal visibility. This project is developing a highly-configurable open-source system-on-chip (SoC) platform that provides both software and hardware models, by using the CHISEL language. CHISEL offers both behavioral (software) simulation and hardware emulation from the same code base. This infrastructure will be freely available for research as well as manufacturing. Currently this project focuses on the on-chip network, but will later expand to the complete SoC to include cores, memory controllers, and other IP blocks.

- Scalability of hardware cache coherency in chip multiprocessors with thousands of cores: Still in the early stages, the question this project aims to answer is what form of cache coherency is prudent for future many-core chip multiprocessors. If no current hardware, software or hybrid scheme is optimal, we will research the bottlenecks and cost tradeoffs, and develop new protocols. The project on rapid

evaluation of future large-scale chips (OpenSoC) will provide a solid infrastructure to provide concrete and detailed data to enable us to complete this study.

- While at Stanford University:

  - Channel reservation protocol for over-subscribed channels and destinations: The channel reservation protocol (CRP) focuses on system-wide networks, such as supercomputer networks. Such networks are often oversubscribed. To make matters worse, workloads in such networks can overstress destinations, creating hotspots. The most popular congestion control technique in today's networks is explicit congestion notification (ECN). However, ECN only acts after congestion is formed and thus reacts slowly to changes in the traffic pattern. Furthermore, ECN's configuration parameters are overly sensitive to the traffic pattern and network design. To mitigate these problems, CRP prevents congestion from ever occurring by allowing sources to reserve bandwidth in multiple resources with a single request, but avoids idling of resources compared to circuit switching. Furthermore, CRP prevents congestion from occurring in traffic patterns where ECN is ineffective, such as congestion formed in network channels by short-lived flows generated by a large combination of source-destination pairs. This way, a benign flow is not affected by tree saturation caused by an adversarial flow. CRP is an extension to the speculative reservation protocol (SRP) [7] because SRP does not reserve bandwidth in network channels.

  - Packet chaining: efficient single-cycle allocation for on-chip networks: Packet chaining improves the matching efficiency of separable allocators without extending cycle time. Packet chaining chains short packets together, to reuse connections of departing packets. This allows an allocator to build up an efficient matching over a number of cycles like incremental allocation, but not limited by packet length. Packet chaining uses a separate allocator which locates packets eligible to reuse connections of departing packets. Packet chaining provides considerable performance gains and enables separable allocators to provide performance comparable to or higher than more expensive and slower allocators. This is important because short packets often dominate chip multiprocessor traffic, and many networks are constrained by cycle time or cannot afford more complex allocators.

  - Evaluating bufferless flow control for on-chip networks: In this work I investigate bufferless flow control and compare it against an optimized buffered on-chip network. As shown, the allocator required by bufferless flow control contains a long serial path because all flits must be sent to an output port without any cycle delay. Moreover, the lack of backpressure means that very long buffers or additional mechanisms are required to prevent deadlocks or packet drops at the endpoints. Also, buffers can be implemented as custom SRAMs which have considerably reduced costs, especially occupied area. Furthermore, this work shows that with buffer bypassing, bufferless flow control provides only 1.5% energy savings at best but still has the side-effects of bufferless flow control outlined in this work. This paper was a best paper award candidate and was very well perceived and cited. It addressed an important debate present at the on-chip network community at the time.

  - An analysis of on-chip networks for large-scale chip multiprocessors: This work is an early attempt to co-design an on-chip network with the cache hierarchy in a chip multiprocessor. This work compares different network topologies by simulating various networks and system parameters, and analyzing the implications in applications and cache parameters. It also uses a cost model to

estimate the network's area and power costs relative to the system. As part of this study, I examine the effect some system-level parameters, such as L2 cache sharing, size and associativity have on network usage and performance. This work also evaluates a hierarchical network design where a small but variable number of processing cores and L2 cache slices are grouped into one tile with a local interconnect, and tiles are connected with a global interconnect.

- Elastic buffer flow-control for on-chip networks: This is my longest-lasting project and is a significant part of my PhD thesis. This work proposes elastic buffer flow control which uses the on-chip channel pipeline flip-flops for buffering instead of router buffers by adding a control logic block to drive the master and slave latch enable inputs independently. This increases the performance over cost efficiency but also significantly simplifies the network thus reducing cycle time. Also, I have designed the two-cycle router, which minimizes cycle time, and the single-cycle router which minimizes latency but is still clocked faster than a virtual channel router. Elastic buffer networks use physical channels instead of virtual channels for traffic separation. I also propose a mechanism to sense congestion because credits are no longer used. Elastic buffers enable networks to remove buffers but maintain buffering.

- Adaptive backpressure: Efficient buffer management for on-chip networks: This work proposes a novel scheme that improves the utilization of dynamically managed router input buffers by continuously adjusting the stiffness of the flow control feedback loop in response to observed traffic conditions. This minimizes the amount of buffer space occupied unproductively by stalled packets and leads to more efficient distribution of buffer space as well as improves isolation between multiple concurrently executing workloads with differing performance characteristics.

- Network Congestion Avoidance through Speculative Reservation: The speculative reservation protocol (SRP) prevents tree saturation from hotspot destinations by having sources reserve bandwidth from their destinations in advance of packet submission. This way, benign flows are not affected by other flows causing the hotspot. SRP uses speculative packets which are sent without waiting for a reservation, but are dropped in case of contention.

- While at ICS-FORTH and the University of Crete:

  - Master's thesis: My MSc thesis proposes a technique to bypass the router pipeline entirely and provide router latencies of potentially less than a cycle. Routers have a set of tri-state drivers or multiplexers, which are pre-set according to pre-configured paths. These paths are run-time re-configurable. Packets traversing those pre-configured paths skip the router pipeline, even before the routing logic computes their desired output. Therefore, this technique essentially performs speculative routing. Copies of packets are kept in router buffers because packets that are mis-routed are dropped at the end of the pre-configured path and re-transmitted at the router where the wrong routing prediction was made. This work provides a solution to reducing on-chip network latency, because that is often the constraint in under-utilized networks.

  - Switch network interface card: I developed a fully-functional PCI-X network interface card to enable communication between a PC and a switching fabric implemented in FPGAs. The design was placed and routed into a Xilinx FPGA.

- Undergraduate degree thesis: I developed a bridge module between the wishbone and PCI interfaces. It also implemented basic flow control and clock frequency domain crossing. A company affiliated with ICS-FORTH then used this bridge module.

COMMUNITY SERVICE

- Served as reviewer for multiple papers for each of ISCA, MICRO, HPCA, SC, IEEE transactions on parallel and distributed systems, IEEE Transactions on Computers (special issue on on-chip networks), ACM Transactions on Computer Architecture and Code Optimization, IEEE Transactions on Communications, IEEE computer architecture letters, IEEE communication letters and Elsevier microprocessors and microsystems journal.
- Program committee member of SC 2016, ICCD 2016-2013, INFOCOMP 2012, and INFOCOMP 2011.

LEADERSHIP EXPERIENCE

- Lecturer at both Stanford University and UC Berkeley.
- Organized a panel for SC15 and organizing another for SC16.
- Organized a BoF session at SC14 and a tutorial at NOCS 2014.
- Co-chair of the Graduate Student Programming Board (GSPB), part of the Graduate Life Office. As a co-chair, I managed the group's annual budget of $43,000 and led a group of approximately fifteen students in organizing events ranging from a few participants all the way to a 200-300.
- FAA certified flight instructor and commercial pilot.
- Civil air patrol member at the grade of captain. Also an instructor, check pilot, and emergency services pilot.
- Community associate for my graduate residence, in charge of social events, emergency response and conflict resolution.
- Tournament official for a strategy card game. This has enabled me to develop leadership skills since I have been the head judge in tournaments of 650 competitors with 35 other officials as my staff. It also allowed me to develop mentoring skills as part of training and certifying new tournament officials.