

Alexandros DAGLIS

ScD, Computer Science
Assistant Professor in Computer Science, Georgia Institute of Technology

Website: www.cc.gatech.edu/~adaglis E-mail: alexandros.daglis@cc.gatech.edu

RESEARCH INTERESTS

Computer architecture, datacenter systems, hardware-software co-design

EDUCATION

- 2012–2018 Doctor of Science degree in COMPUTER SCIENCE
École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland
Thesis: “Network-compute co-design for distributed in-memory computing”
Advisors: Prof. Babak FALSAFI & Prof. Edouard BUGNION
- 2007–2012 Diploma in ELECTRICAL & COMPUTER ENGINEERING (5-year degree)
National Technical University of Athens (NTUA), Athens, Greece
Thesis: “A study of a dynamic placement policy in a NUCA cache”
Advisor: Prof. Nectarios KOZIRIS

AWARDS & HONORS

- Fall 2023 NSF CAREER Award.
- Fall 2022 Georgia Tech College of Computing Outstanding Junior Faculty Teaching Award.
- Fall 2022 Georgia Tech Student Recognition of Excellence in Teaching.
- Fall 2021 Georgia Tech Student Recognition of Excellence in Teaching.
- 2019 Google Faculty Research Award.
- 2019 ACM SIGARCH/IEEE CS TCCA Outstanding Dissertation – Honorable Mention for “contributions to network-centric server architecture for in-memory data-center services”.
- 2018 Nominated for an ACM Doctoral Dissertation Award by EPFL’s School of Computer and Communication Sciences.
- 2018 EPFL remarkable thesis distinction. Awarded to top 8% of PhD dissertations.
- 2013–2014 Microsoft Research Fellowship.
- 2012–2013 EPFL Computer Science Fellowship.

RESEARCH FUNDING

- 2023 NSF CAREER—*Architecting Datacenters for Optimized Tail Latency at Scale*
- 2023 Samsung Semiconductor Inc—*Rethinking Server Memory Hierarchy in the Era of CXL*.
- 2022 IARPA AGILE—*Scalable Locally-Centric Graph Analytics System*.

- 2022 NSF award 2212098—*Collaborative Research: CNS Core: Medium: High-performance Network Stacks for the Edge.*
- 2022 Intel Corporation—*TRIM: Architecting Systems for Terascale in-Memory Applications.*
- 2021 Samsung Semiconductor Inc—*Disaggregated Memory Architectures for Scalable Distributed Training.*
- 2020 NSF award 2006602—*SHF: CNS Core: Small: Server architecture optimizations for microsecond-scale RPCs.*

EMPLOYMENT

JAN. '19–PRESENT	Assistant Professor at School of Computer Science, College of Computing, Georgia Institute of Technology
SEP. '12–SEP. '18	Research Assistant at EPFL, Lausanne, Switzerland Research on Datacenter Technologies and In-Memory Rack-scale Computing <ul style="list-style-type: none"> • Lead architect of the Scale-Out NUMA project.
JAN.–MAY '15	Systems Research Intern at HP LABS, Palo Alto, USA Software and hardware research for HP's <i>The Machine</i>

TEACHING

- Systems and Networks (cs2200) Spring 2022, Spring 2023
- Topics on Datacenter Design (cs8803) Fall 2019, Spring 2020, Spring 2021
- Advanced Computer Architecture Fall 2020, Fall 2021, Fall 2022, Fall 2023
(cs4290/6290 – ece4100/6100)

CONFERENCE PUBLICATIONS

- [1] **Safety Hints for HTM Capacity Abort Mitigation.** Anirudh Jain, Divya Kiran Kadiyala, Alexandros Daglis. In Proceedings of the *29th IEEE International Symposium on High-Performance Computer Architecture (HPCA-29)*, Montreal, Canada, 2023.
- [2] **Turbo: SmartNIC-enabled Dynamic Load Balancing of μ S-scale RPCs.** Hamed Seydroudbari, Srikar Vanavasam, Alexandros Daglis. In Proceedings of the *29th IEEE International Symposium on High-Performance Computer Architecture (HPCA-29)*, Montreal, Canada, 2023.
- [3] **Cooperative Concurrency Control for Write-Intensive Key-Value Workloads.** Mark Sutherland, Babak Falsafi, Alexandros Daglis. In Proceedings of the *28th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XXVIII)*, Vancouver, Canada, 2023.

- [4] **Patching up Network Data Leaks with Sweeper.** Marina Vemmou, Albert Cho, Alexandros Daglis. In Proceedings of the *55th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Chicago, IL, USA, 2022.
- [5] **OneEdge: An Efficient Control Plane for Geo-Distributed Infrastructures.** Enrique Saurez, Harshit Gupta, Alexandros Daglis, Umakishore Ramachandran. In Proceedings of the *ACM Symposium on Cloud Computing (SoCC)*, Seattle, WA, USA, 2021.
- [6] **COSPlay: Leveraging Task-Level Parallelism for High-Throughput Synchronous Persistence.** Marina Vemmou and Alexandros Daglis. In Proceedings of the *54th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Worldwide Event, 2021.
- [7] **Cerebros: Evading the RPC Tax in Datacenters.** Arash Pourhabibi, Mark Sutherland, Alexandros Daglis, Babak Falsafi. In Proceedings of the *54th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Worldwide Event, 2021.
- [8] **The NeBuLa RPC-Optimized Architecture.** Mark Sutherland, Siddharth Gupta, Babak Falsafi, Virendra Marathe, Dionisios Pnevmatikatos, Alexandros Daglis. In Proceedings of the *47th International Symposium on Computer Architecture (ISCA)*, Worldwide Event, 2020.
- [9] **Distributed Logless Atomic Durability with Persistent Memory.** Siddharth Gupta, Alexandros Daglis, Babak Falsafi. In Proceedings of the *52nd International Symposium on Microarchitecture (MICRO)*, Columbus, OH, USA, 2019.
- [10] **RPCValet: NI-Driven Tail-Aware Balancing of μ s-scale RPCs.** Alexandros Daglis, Mark Sutherland, Babak Falsafi. In Proceedings of the *24th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XXIV)*, Providence, RI, USA, 2019.
- [11] **Design Guidelines for High-Performance SCM Hierarchies.** Dmitrii Ustiugov, Alexandros Daglis, Javier Picorel, Mark Sutherland, Edouard Bugnion, Babak Falsafi, Dionisios Pnevmatikatos. In Proceedings of the *4th Annual International Symposium on Memory Systems (MEMSYS)*, Washington DC, USA, 2018.
- [12] **The Mondrian Data Engine.** Mario Drumond, Alexandros Daglis, Nooshin Mirzadeh, Dmitrii Ustiugov, Javier Picorel, Babak Falsafi, Boris Grot, Dionisios Pnevmatikatos. In Proceedings of the *44th International Symposium on Computer Architecture (ISCA)*, Toronto, ON, Canada, 2017.
- [13] **The Case for RackOut: Scalable Data Serving Using Rack-Scale Systems.** Stanko Novakovic, Alexandros Daglis, Edouard Bugnion, Babak Falsafi, Boris Grot. In Proceedings of the *ACM Symposium on Cloud Computing (SoCC)*, Santa Clara, CA, USA, 2016.
- [14] **SABRes: Atomic Object Reads for In-Memory Rack-Scale Computing.** Alexandros Daglis, Dmitrii Ustiugov, Stanko Novakovic, Edouard Bugnion, Babak Falsafi, Boris Grot. In Proceedings of the *49th International Symposium on Microarchitecture (MICRO)*, Taipei, Taiwan, 2016.
- [15] **An Analysis of Load Imbalance in Scale-out Data Serving.** Stanko Novakovic, Alexandros Daglis, Edouard Bugnion, Babak Falsafi, Boris Grot. *ACM SIGMETRICS* (Short paper), Antibes Juan-Les-Pins, France, 2016.
- [16] **Manycore Network Interfaces for In-Memory Rack-Scale Computing.** Alexandros Daglis, Stanko Novakovic, Edouard Bugnion, Babak Falsafi, Boris Grot. In Proceedings of the *42nd International Symposium on Computer Architecture (ISCA)*, Portland, OR, USA, 2015.

- [17] **Scale-Out NUMA.** Stanko Novakovic, Alexandros Daglis, Edouard Bugnion, Babak Falsafi, Boris Grot. In *Proceedings of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XIX)*, Salt Lake City, UT, USA, 2014.

JOURNAL PUBLICATIONS

- [18] **IDIO: Orchestrating Inbound Network Data on Server Processors.** Mohammad Alian, Jongmin Shin, Ki-Dong Kang, Ren Wang, Alexandros Daglis, Daehoon Kim, Nam Sung Kim. *IEEE Computer Architecture Letters*, 2020.
- [19] **Exploiting Errors for Efficiency.** Phillip Stanley-Marbell, Armin Alaghi, Michael Carbin, Eva Darulova, Lara Dolecek, Andreas Gerstlauer, Ghayoor Gillani, Djordje Jevdjic, Thierry Moreau, Mattia Cacciotti, Alexandros Daglis, Natalie Enright Jerger, Babak Falsafi, Sasa Misailovic, Adrian Sampson, Damien Zufferey. *ACM Computing Surveys*, 2020.
- [20] **Mitigating Load Imbalance in Distributed Data Serving with Rack-Scale Memory Pooling.** Stanko Novakovic, Alexandros Daglis, Dmitrii Ustiugov, Edouard Bugnion, Babak Falsafi, Boris Grot. *ACM Transactions on Computer Systems (TOCS)*, 2019.
- [21] **Algorithm/Architecture Co-Design for Near-Memory Processing.** Mario Drumond, Alexandros Daglis, Nooshin Mirzadeh, Dmitrii Ustiugov, Javier Picorel, Babak Falsafi, Boris Grot, Dionisios Pnevmatikatos. *ACM SIGOPS Operating Systems Review*, 2018.

WORKSHOP PUBLICATIONS, ARTICLES & ARXIV PUBLICATIONS

- [22] **Filtering Wasteful Vertex Visits in Breadth-First Search.** Prachatos Mitra, Alexandros Daglis. In *Proceedings of the SC '23 Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis (SC-W)*, 2023.
- [23] **A Case for CXL-Centric Server Processors.** Albert Cho, Anish Saxena, Moinuddin Qureshi, Alexandros Daglis. <https://arxiv.org/abs/2305.05033>, 2023.
- [24] **Exploring Memory Expansion Designs for Training Mixture-of-Experts Models.** Taekyung Heo, Saeed Rashidi, Changhai Man, Divya Kiran Kadiyala, William Won, Sudarshan Srinivasan, Midhilesh Elavazhagan, Madhu Kumar, Alexandros Daglis, Tushar Krishna. *Workshop on Hot Topics in System Infrastructure (HotInfra)*, 2023.
- [25] **NFSlicer: Data Movement Optimization for Shallow Network Functions.** Anirudh Sarma, Hamed Seyedroudbari, Harshit Gupta, Umakishore Ramachandran, Alexandros Daglis. <https://arxiv.org/abs/2203.02585>, March 2022.
- [26] **eCloud: A Vision for the Evolution of the Edge-Cloud Continuum.** Joy Arulraj, Abhijit Chatterjee, Alexandros Daglis, Ashutosh Dhekne, and Umakishore Ramachandran. *IEEE Computer, Special Issue on Computing for Autonomy: Latency, Power, Resilience*, May 2021.
- [27] **Unleashing the Full Potential of Persistent Memory with Logless Atomic Durability.** Siddharth Gupta, Alexandros Daglis, Babak Falsafi. *11th Annual Non-Volatile Memories Workshop*, 2020.

INVITED TALKS

A Case for CXL-Centric Server Processors	
• Microsoft Azure Hardware Architecture AHA Learning Series	October 2023
Architecting Systems for Terascale in-Memory Applications	
• Intel Center for Transformative Server Architectures (TSA) Annual Workshop	September 2023
Rethinking the Network-Compute Interface in the era of extreme software decomposition	
• University of Washington	March 2023
A decade of Scale-Out NUMA: Impact on my PhD, early career, and research community	
• Babak Falsafi's 25-year Career Anniversary Workshop	February 2023
Efficient Large-Scale Architectures for Distributed Training Using Memory Expansion Techniques	
• Samsung Memory Solutions Lab	May 2022
Optimizing the "Last Mile" with Network-Compute Co-Design	
• Center for Research into Novel Computing Hierarchies (CRNCH) Summit	Jan. 2021
Network/Architecture CoDesign	
• <i>Happy Hour with Architects</i> online series, hosted by Samira Khan	May 2020
RPCValet: NI-Driven Tail-Aware Balancing of μs-scale RPCs	
• 24th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-XXIV), Providence, RI, USA (conference presentation)	Apr. 2019
Network-Centric Computing for Online Services	
• CS seminar at Stony Brook University	Sep. 2019
• ECE seminar at the National Technical University of Athens (NTUA)	Sep. 2018
• CS research seminar at Imperial College London	Apr. 2018
• CIS seminar at the University of Pennsylvania	Mar. 2018
• EE departmental seminar at Princeton University	Mar. 2018
• CS research seminar at Rutgers University	Mar. 2018
• CS research seminar at UCLA	Feb. 2018
• SCS seminar at Georgia Tech	Feb. 2018
SABRes: Atomic Object Reads for In-Memory Rack-Scale Computing	
• 49th International Symposium on Microarchitecture (MICRO), Taipei, Taiwan (conference presentation)	Oct. 2016
Chip Design for In-Memory Rack-Scale Computing	
• EcoCloud Annual Event, Lausanne, Switzerland	May 2016
Manycore Network Interfaces for In-Memory Rack-Scale Computing	

- Computing Systems Research Day at the National Technical University of Athens (NTUA), Athens, Greece Jan. 2016
 - 42nd International Symposium in Computer Architecture (ISCA), Portland, USA (conference presentation) Jun. 2015
- Hybrid Coherence for “The Machine”**
- HP Labs, Palo Alto, USA May 2015

PROFESSIONAL SERVICE

PC MEMBER / REVIEWER:

- Conference Program Committee: ASPLOS 2024, ASPLOS 2023, ISCA 2022 Industry Track, HPCA 2022, ASPLOS 2021, USENIX ATC 2021, USENIX ATC 2020, ISPASS 2020, CLUSTER 2019
- Conference External Review Committee: ISCA 2023, ISCA 2021, ISCA 2020, MICRO 2020
- Journals: CAL 2021, CAL 2020, IEEE Transactions on Computers 2018/2023, TACO 2019
- Workshops: HotInfra 2023, YArch 2022, WAX 2019, WORD 2019
- Other: MICRO Student Research Competition 2019

ORGANIZATIONAL ROLES:

- MICRO Artifact Evaluation Co-Chair, 2022
- ASPLOS-XXVII Travel Grants Chair, 2022
- [Young Architect \(YArch\)](#) Workshop co-organizer, 2021 (at ASPLOS-XXVI)
- [Young Architect \(YArch\)](#) Workshop co-organizer and PC chair, 2020 (at ASPLOS-XXV)
- [Rising Stars in Computer Architecture \(RISC-A\)](#) Workshop co-organizer, 2019

RESEARCH PROPOSAL EVALUATIONS:

- Israel Science Foundation reviewer, 2023
- NSF proposal review panelist, 2019
- NASA reviewer, 2019

MENTORING

PhD	Prachatos Mitra	2023-
	Albert Cho	2021-
	Marina Vemmou	2019-
	Hamed Seyedroudbari	2019-
	Divya Kadiyala	2019-
	Mark Sutherland (w/ Babak Falsafi)	2018–2022
MS (project)	Prashant Ramnani	2023-
	Rachana Aithal	2023-
	Dhruva Devasthale	2023
	Manish Manchali	2023

	Sibi Sudhakar	2023
	Jayant Tandon	2023
	Aditya Rohan	2022
	Balaji Ravikumar	2022
	Nagasayee Guduru	2022
	Harigovind Anil	2022
	Dhruva Barfiwala	2022
	Arvind Sivasankar	2022
	Sriyash Caculo	2021–2022
BS (thesis)	Srikar Vanavasam	2022–2023

PATENTS

- [1] *Atomic Object Reads for In-Memory Rack-Scale Computing* | US Patent 10,929,174, 2021
 - With Boris Grot and Babak Falsafi.
- [2] *Scale-Out Non-Uniform Memory Access* | US Patent 9,734,063, 2017
 - With Stanko Novakovic, Boris Grot, Edouard Bugnion, and Babak Falsafi.
 - Licensed by a major IT vendor in 2016.

LANGUAGES

GREEK, ENGLISH, GERMAN (limited working proficiency), FRENCH (elementary proficiency).