

BORIS GROT

CONTACT INFORMATION

University of Edinburgh
School of Informatics
Informatics Forum, Rm 2.04A
10 Crichton Street
Edinburgh, EH8 9AB
United Kingdom

Web: <http://homepages.inf.ed.ac.uk/bgrot>
E-mail: boris.grot@ed.ac.uk
Tel: +44 131 650 5123

RESEARCH INTERESTS

- **Broad:** Computer architecture, systems, networks.
- **Current focus:** System architectures for datacenters, architectural support for large-scale data analytics, rack-scale systems, systems with quality-of-service guarantees, interconnection networks.

EDUCATION

University of Texas at Austin Austin, TX
Ph.D., Computer Science Aug 2011

- Thesis: “Network-on-Chip Architectures for Scalability and Service Guarantees”
- Advisor: Prof. Stephen W. Keckler

University of California, Los Angeles Los Angeles, CA
M.S., Electrical Engineering Dec 2005

- Thesis: “Precise Flow Tracking in High-speed Networks: Memory Performance Analysis and Techniques for Performance Enhancement”
- Advisor: Prof. William Mangione-Smith

Pennsylvania State University State College, PA
B.S. (with Honors), Computer Engineering May 2000

PROFESSIONAL EXPERIENCE

- Reader (US title: Associate Professor), School of Informatics, University of Edinburgh (Aug 2018-Present)
- Lecturer (US title: Assistant Professor), School of Informatics, University of Edinburgh (Jan 2014-Jul 2018)
- Post-doctoral researcher, *EPFL* (2011-2013)
- Consultant, *NVIDIA Research* (2010-2011)
- Intern, *Microsoft Research* (Summer 2008)
- Hardware Engineering Intern, *Intel* (Summer 2005)
- Hardware Engineer, *Calix Networks* (2001-2002)
- Verification Engineer, *Lucent Technologies* (2000-2001)

AWARDS AND RECOGNITION

- ASPLOS Distinguished Artifact Award, 2021
- IEEE Micro Top Picks in Computer Architecture - Honorable Mention, 2020

- PPOPP Best Paper Nominee, 2020
- HPCA Best Paper Award, 2019
- ARM Research Summit Best Poster Award, 2018, 2019
- MICRO Hall of Fame inductee, 2016
- Google Faculty Award recipient: 2019, 2014
- MICRO Best Paper Runner-Up, 2013
- IEEE Micro Top Pick in Computer Architecture, 2012
- University of Edinburgh Teaching Award nominee, 2015-2019

REFEREED CONFERENCE PUBLICATIONS

- A. Katsarakis, Y. Ma, Z. Tan, A. Bainbridge, M. Balkwill, A. Dragojevic, B. Grot, B. Radunovic, Y. Zhang. Zeus: Locality-aware Distributed Transactions. In *European Conference on Computer Systems (EuroSys)*, 2021.
- D. Ustiugov, P. Petrov, M. Kogias, E. Bugnion, B. Grot. Benchmarking, Analysis, and Optimization of Serverless Function Snapshots. In *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2021. **Distinguished Artifact Award**
- A. Margaritov, D. Ustiugov, A. Shahab, B. Grot. PTEMagnet: Fine-grained Physical Memory Reservation for Faster Page Walks in Public Clouds. In *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2021.
- A. Shahab and B. Grot. Population-based Evolutionary Distributed SGD [Poster]. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, 2020.
- M. R. S. Katebzadeh, P. Costa, B. Grot. Evaluation of an InfiniBand Switch: Choose Latency or Bandwidth, but Not Both. In *International Symposium on Performance Analysis of Systems and Software (ISPASS)*, 2020.
- A. Katsarakis, V. Gavrielatos, A. Joshi, M.R.S. Katebzadeh, B. Grot, V. Nagarajan, A. Dragojevic. Hermes: Fast, Fault-Tolerant and Linearizable Data Replication. In *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2020. **IEEE Micro Top Picks in Computer Architecture – Honorable Mention.**
- V. Gavrielatos, A. Katsarakis, V. Nagarajan, B. Grot, A. Joshi. Kite: efficient and available release consistency for the datacenter. In *Symposium on Principles and Practice of Parallel Programming (PPOPP)*, 2020. **Best Paper Nominee.**
- P. Faldu, J. Diamond, B. Grot. Domain-Specialized Cache Management for Graph Analytics. In *International Symposium on High-Performance Computer Architecture (HPCA)*, 2020.
- P. Faldu, J. Diamond, B. Grot. A Closer Look at Lightweight Graph Reordering. In *International Symposium on Workload Characterization (IISWC)*, 2019.
- A. Margaritov, D. Ustiugov, E. Bugnion, B. Grot. Prefetched Address Translation. In *52nd International Symposium on Microarchitecture (MICRO)*, 2019.
- A. Margaritov, S. Gupta, R. Gonzalez-Alberquilla, B. Grot. Stretch: Balancing QoS and Throughput for Colocated Server Workloads on SMT Cores. In *25th International Symposium on High-Performance Computer Architecture (HPCA)*, 2019. **Best Paper Award.**
- A. Shahab, M. Zhu, A Margaritov, B. Grot. Farewell My Shared LLC! A Case for Private Die-Stacked DRAM Caches for Servers. In *51st International Symposium on Microarchitecture (MICRO)*, 2018.
- V. Gavrielatos, A. Katsarakis, A. Joshi, N. Oswald, B. Grot, V. Nagarajan. Scale-Out ccNUMA:

Exploiting Skew with Strongly Consistent Caching. In *European Conference on Computer Systems (EuroSys)*, 2018.

- R. Kumar, B. Grot, V. Nagarajan. Blasting Through The Front-End Bottleneck With Shotgun. In *23rd International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2018.
- P. Faldu and B. Grot. Leeway: Addressing Variability in Dead-Block Prediction for Last-Level Caches. In *26th International Conference on Parallel Architectures and Compilation Techniques (PACT)*, 2017.
- M. Drumond, A. Daglis, D. Ustiugov, N. Mirzadeh, J. Picorel, B. Falsafi, B. Grot, D. Pnevmatikatos. The Mondrian Data Engine. In *44th International Symposium on Computer Architecture (ISCA)*, 2017.
- R. Kumar, C. Huang, B. Grot, V. Nagarajan. Boomerang: a Metadata-Free Architecture for Control Flow Delivery. In *23rd International Symposium on High Performance Computer Architecture (HPCA)*, 2017.
- C. Huang, R. Kumar, M. Elver, B. Grot, V. Nagarajan. C3D: Mitigating NUMA Effects via Coherent DRAM Caches. In *49th International Symposium on Microarchitecture (MICRO)*, 2016.
- A. Daglis, D. Ustiugov, S. Novakovic, E. Bugnion, B. Falsafi, B. Grot. SABRes: Fast Atomic Remote Object Reads for Rack-Scale In-Memory Computing. In *49th International Symposium on Microarchitecture (MICRO)*, 2016.
- S. Novakovic, A. Daglis, E. Bugnion, B. Falsafi, B. Grot. The Case for RackOut: Scalable Data Serving Using Rack-Scale Systems. In *7th ACM Symposium on Cloud Computing (SOCC)*, 2016.
- O. Kocberber, B. Grot, and B. Falsafi. Asynchronous Memory Access Chaining. In *42nd International Conference on Very Large Data Bases (VLDB)*, 2016.
- C. Kaynak, B. Grot, and B. Falsafi. Confluence: Unified Instruction Supply for Scale-Out Servers. In *48th International Symposium on Microarchitecture (MICRO)*, 2015.
- A. Daglis, S. Novakovic, E. Bugnion, B. Falsafi, B. Grot. Manycore Network Interfaces for In-Memory Rack-Scale Computing. In *International Symposium on Computer Architecture (ISCA)*, 2015.
- S. Volos, J. Picorel, B. Falsafi, B. Grot. BuMP: Bulk Memory Page Access Prediction and Streaming. In *International Symposium on Microarchitecture (MICRO)*, 2014.
- S. Fytraki, E. Vlachos, O. Kocberber, B. Falsafi, B. Grot. FADE: A Programmable Filtering Accelerator for Instruction-Grain Monitoring. In *International Symposium on High Performance Computer Architecture (HPCA)*, 2014.
- S. Novakovic, A. Daglis, E. Bugnion, B. Falsafi, B. Grot. Scale-Out NUMA. In *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2014.
- O. Kocberber, B. Grot, J. Picorel, B. Falsafi, K. Lim, P. Ranganathan. Meet the Walkers: Accelerating Index Traversals for In-Memory Databases. In *International Symposium on Microarchitecture (MICRO)*, 2013. **Best Paper Runner-Up**.
- C. Kaynak, B. Grot, B. Falsafi. SHIFT: Shared History Instruction Fetch for Lean-Core Server Processors. In *International Symposium on Microarchitecture (MICRO)*, 2013.
- P. Lotfi-Kamran, B. Grot, B. Falsafi. NOC-Out: Microarchitecting a Scale-Out Processor. In *International Symposium on Microarchitecture (MICRO)*, 2012.
- P. Lotfi-Kamran, B. Grot, M. Ferdman, S. Volos, O. Kocberber, J. Picorel, A. Adileh, D. Jevdjic, S. Idgunji, E. Ozer, B. Falsafi. Scale-Out Processors. In *International Symposium on Computer Architecture (ISCA)*, 2012.
- S. Volos, C. Seiculescu, B. Grot, N. Khosro Pour, B. Falsafi, G. De Micheli. CCNoC: Specializing On-Chip Interconnects for Energy Efficiency in Cache-Coherent Servers. In *International Symposium on*

Networks-on-Chip (NOCS), 2012.

- B. Grot, J. Hestness, S. W. Keckler, O. Mutlu. Kilo-NOC: A Heterogeneous Network-on-Chip Architecture for Scalability and Service Guarantees. In *International Symposium on Computer Architecture (ISCA)*, 2011. **IEEE Micro Top Pick in Computer Architecture**.
- H. Kim, P. Ghoshal, B. Grot, P. Gratz, D. Jimenez. Reducing Network-on-Chip Energy Consumption through Spatial Locality Speculation. In *International Symposium on Networks-on-Chips (NOCS)*, 2011.
- B. Grot, S. W. Keckler, O. Mutlu. Preemptive Virtual Clock: A Flexible, Efficient, and Cost-effective QOS Scheme for Networks-on-a-Chip. In *International Symposium on Microarchitecture (MICRO)*, 2009.
- B. Grot, J. Hestness, S. W. Keckler, O. Mutlu. Express Cube Topologies for On-Chip Interconnects. In *International Symposium on Computer Architecture (HPCA)*, 2009.
- P. Gratz, B. Grot, S. W. Keckler. Regional Congestion Awareness for Load Balance in Networks-on-Chip. In *International Symposium on High-Performance Computer Architectures (HPCA)*, 2008.

JOURNAL PUBLICATIONS

- S. Novakovic, A. Daglis, D. Ustiugov, E. Bugnion, B. Falsafi, B. Grot. Mitigating Load Imbalance in Distributed Data Serving with Rack-Scale Memory Pooling. In *ACM Transactions on Computer Systems (TOCS)*, Vol. 36, No. 2, 2019.
- M. Drumond, A. Daglis, D. Ustiugov, N. Mirzadeh, J. Picorel, B. Falsafi, B. Grot, D. Pnevmatikatos. Algorithm/Architecture Co-Design for Near-Memory Processing. In *ACM SIGOPS Operating Systems Review*, volume 52, issue 1. 2018.
- S. Volos, D. Jevdjic, B. Falsafi, B. Grot. Fat Caches for Scale-Out Servers. In *IEEE Micro*, volume 37, issue 2. 2017.
- B. Grot, D. Hardy, P. Lotfi-Kamran, C. Nicopoulos, Y. Sazeides, B. Falsafi. Optimizing Datacenter TCO with Scale-Out Processors. In *IEEE Micro, Special Issue on Energy-Aware Computing*, volume 32, issue 5, 2012.
- B. Grot, J. Hestness, S. W. Keckler and O. Mutlu. A QoS-Enabled On-Die Interconnect Fabric for Kilo-Node Chips. In *IEEE Micro, Top Picks from 2011 Computer Architecture Conferences*, volume 32, issue 3, 2012. (original in ISCA 2011)
- H. Kim, B. Grot, P. V. Gratz, D. A. Jimenez. Spatial Locality Speculation to Reduce Energy in Chip-Multiprocessor Networks-on-Chip, In *IEEE Transactions on Computers, Special Issue on NOCS*, 2012.

REFEREED WORKSHOP PUBLICATIONS

- D. Ustiugov, P. Petrov, M.R.S. Katebzadeh, B. Grot. Bankrupt Covert Channel: Turning Network Predictability into Vulnerability. In *USENIX Workshop on Offensive Technologies (WOOT), co-located with USENIX Security*, 2020.
- Artemiy Margaritov, Dmitri Ustiugov, Edouard Bugnion, and Boris Grot. Virtual Address Translation via Learned Page Table Indexes. In *Workshop on ML for Systems at NIPS*, 2018.
- P. Faldu and B. Grot. Leeway: Highly Adaptive Cache Management. In *2nd Cache Replacement Competition (CRC2)*, 2017.
- P. Faldu and B. Grot. LLC Dead Block Prediction Considered Not Useful. In *13th Workshop on Duplicating, Deconstructing and Debunking (WDDD)*, 2016.

- N. Mirzadeh, O. Kocberber, B. Falsafi, B. Grot. Sort vs. Hash Join Revisited for Near-Data Execution. In *Fifth Workshop on Architectures and Systems for Big Data (ASBD)*, 2015.
- J. Hestness, B. Grot, S. W. Keckler. Netrace: Dependency-Driven Trace-Based Network-on-Chip Simulation. In *Workshop on Network on Chip Architectures (NOCARC)*, 2010.
- B. Grot, S. W. Keckler and O. Mutlu. Topology-aware Quality-of-Service Support in Highly Integrated Chip Multiprocessors. In *Workshop on the Interaction between Operating Systems and Computer Architecture (WIOSCA)*, 2010.
- S. Prabhu, B. Grot, P. V. Gratz and J. Hu. Ocim_tsim: a DVFS-aware simulator for NoC based platforms. In *Workshop on SoC Architecture, Accelerators and Workloads (SAW)*, 2010.
- K. C. Hale, B. Grot, S. W. Keckler. Segment Gating for Static Energy Reduction in Networks-On-Chip. In *Workshop on Network on Chip Architectures (NOCARC)*, 2009.
- B. Grot and S. W. Keckler. Scalable On-Chip Interconnect Topologies. In *Workshop on Chip Multiprocessor Memory Systems and Interconnects (CMP-MSI)*, 2008.
- B. Grot and W. Mangione-Smith. Good Memories: Enhancing Memory Performance for Precise Flow Tracking. In *Advanced Networking and Communications Hardware Workshop (ANCHOR)*, 2005.

PROFESSIONAL ACTIVITIES

- Board of Directors, ACM Special Interest Group on Computer Architecture (ACM SIGARCH). Aug 2019 – Present.
- Information Director, ACM Special Interest Group on Computer Architecture (ACM SIGARCH). Sep 2014 – Jul 2019.
- Program Committee Member (excluding External Review Committees): IEEE MICRO Top Picks (2020, 2019, 2018), ISCA (2021, 2018, 2016, 2014), ASPLOS (2020, 2019, 2018, 2015), MICRO (2020, 2016, 2015, 2014, 2013, 2012), HPCA (2022, 2018, 2016, 2015), ATC (2019), SYSTOR (2018), IISWC (2017, 2014, 2012), ISPASS (2016), ICS (2014)
- Journal reviewer: IEEE TOCS, IEEE CAL, IEEE TCAD, IEEE TPDS, ACM TODAES, ACM TACO
- Guest co-editor, IEEE Micro, Special Issue on Near-Memory Processing, Jan-Feb 2016
- Guest co-editor, IEEE Micro, Special Issue on Big Data, Jul-Aug 2014
- Organizer, Workshop on Warehouse-Scale Memory Systems (WAMS), held in conjunction with ASPLOS 2018.
- Organizer, Workshop on Rack-Scale Memory Systems and Models (RAMbO), held in conjunction with HiPEAC 2017
- Organizer and PC co-chair, 1st Workshop on Multicore and Rack-scale Systems (MaRS), held in conjunction with EuroSys 2016
- Organizer and PC co-chair, 5th Workshop on Architectures and Systems for Big Data (ASBD), held in conjunction with ISCA 2015
- Organizer and PC co-chair, Workshop on Near-Data Processing (WoNDP), held in conjunction with MICRO 2014
- Mini-Symposia chair, ParCo 2015
- Publicity chair: EuroSys 2020, ISCA 2018, ASPLOS 2018, LCTES 2017, IISWC 2017
- Workshop/Tutorial chair: MICRO 2020, ASPLOS 2017
- Poster chair, PLDI 2014
- Web chair, IEEE Micro Top Picks 2013
- Workshop PC Member: ASBD (2014), MSPC (2012)

PATENTS

- *Atomic Object Reads for In-Memory Rack-scale Computing*. With A. Daglis and B. Falsafi. US Patent # US 10,929,174 B2. Issued Feb 23, 2021.
- *Unified Prefetching Into Instruction Cache And Branch Target Buffer*. With B. Falsafi and C. Kaynak. US Patent # US 9,996,358 B2. Issued June 12, 2018.
- *Network-on-Chip Using Request and Reply Trees for Low-Latency Processor-Memory Communication*. With B. Falsafi and P. Lotfi-Kamran. US patent #9,703,707. Issued July 11, 2017.
- *Scale-Out Non-Uniform Memory Access*. With S. Novakovic, A. Daglis, E. Bugnion and B. Falsafi. US Patent #20,150,242,324. Aug 27, 2015.
- *Method and Apparatus for Congestion-Aware Routing in a Computer Interconnection Network*. With P. Gratz and S. W. Keckler. US Patent #8,285,900. Issued Oct 9, 2012.
- *Scalable Bus-Based On-Chip Interconnection Networks*. With S. W. Keckler. US Patent #8,307,116. Issued Nov 6, 2012.

FUNDING

- *Google Faculty Research Award*. Accelerating Address Translation via a Learned Page Table Index. \$73,000 (sole PI), 2019
- *DARPA HIVE project (sub-contract with Intel)*. Modeling and Improving Performance of Large-Scale Graph Analytics through Software/Hardware Co-Design. \$77,000 (sole PI), 2018
- *Oracle ERO Award*. Efficient Scale-Up Graph Processing on Future Memory Systems. \$76,000 (sole PI), 2017
- *GCHQ Research Grant*. Networked Server Testbed. £19,500 (sole PI), 2016.
- *Microsoft Research PhD Scholarship*. Rack-Scale Interconnects for Disaggregated Memory (sole PI), 2016
- *Google Faculty Research Award*. Near-Memory Processing for Analytics. \$53,000 (sole PI), 2014.
- *EPSRC Research Grant*. Error-tolerant Stream Processing System Design. £441,000 (co-PI), 2014.
- *CHIST-ERA Research Grant*. Distributed Heterogeneous Vertically-Integrated Energy-Efficient Data Centres. €202,000 (Lead PI); total award: €1.45M (UoE, QUB, Lancaster, EPFL, AMD), 2014.

INVITED TALKS

- “Improving Cache Efficiency of Graph Analytics by Exploiting Graph Structure.” Invited talk. ARM Research Summit. Cambridge UK, Sep 2017.
- “Moore with Less: Specializing Cores for the Cloud”, Invited talk. Technion University. Haifa, Israel. May 2017.
- “Open Source Software in the Datacenter”. Invited talk. HiPEAC Computer Systems Week. Zagreb, April 2017.
- “Lean and Mean Control-Flow Delivery for High-Performance Cores”. Invited talk. Samsung. Austin, USA. February 2017.
- “Kissing the NUMA Bottleneck Goodbye”. Invited talk. ARM Research Summit. Cambridge UK, Sep 2016.
- “Turbo-charging In-Memory Rack-Scale Computing with Scale-Out NUMA”. Invited talk. Microsoft Faculty Summit. Seattle, Jul 2016.
- “Scale-Out NUMA”. Invited talk. Imperial College. London, UK. May 2016.
- “Toward PetaRAM servers with Scale-Out NUMA.” Invited keynote. EuroServer Workshop.

HiPEAC, Amsterdam, Jan 2015.

- “Chip-Level Implications of Rackscale Computing.” Invited talk. Mini-Symposium on Bridging the Gap between Networking and Computing. Cambridge, Dec 2014
- “Future of Architecture and Systems for Big Data Processing.” Invited panelist. ASBD workshop, co-located with ISCA. Minneapolis, Jun 2014.
- “Scale-Out Datacenters: Big Data, Big Servers, Big Trouble”. Invited talk. CUSO Winter School on Data-Centric Systems, Jan 2014.
- “Network-on-Chip for Scale-Out Processors.” Invited keynote. Workshop on Interconnection Network Architectures: On-Chip, Multi-Chip (INA-OCMC). HiPEAC, Berlin, Jan 2013.
- “Toward Sustainable Datacenters.” Invited talk. Swiss National Conference on Supporting Science with Cloud Computing. Bern, Nov 2012.
- “Scale-Out Processors.” Invited talk. Intel European Research & Innovation Conference. Barcelona, Oct 2012.

TEACHING

- *Introduction to Computer Systems*, Univ. of Edinburgh (Fall 2020, 2019, 2018, 2017, 2016, 2015, 2014). Nominated for the **Student Teaching Award** every semester taught.
- *Computer Architecture*, Univ. of Edinburgh (Spring 2018, 2017, 2016, 2015, 2014)
- *Advanced Topics in Datacenter Design*, EPFL (Spring 2013)

STUDENT SUPERVISION

- PhD students (current):
 - Artemiy Margaritov, Univ. of Edinburgh, 2014 – present
 - Amna Shahab, Univ. of Edinburgh, 2015 – present
 - Antonis Katsarakis, Univ. of Edinburgh, 2016 – present
 - Vasileios Gavrielatos, Univ. of Edinburgh, 2016 – present (co-supervised with Dr. Vijay Nagarajan)
 - Mingan Zhu, Univ. of Edinburgh, 2016 – present
 - Siavash Katebzadeh, Univ. of Edinburgh, 2017 – present
 - Dmitrii Ustiugov, Univ. of Edinburgh, 2019 – present
 - David Schall, Univ. of Edinburgh, 2020 – present
- PhD students (graduated):
 - Priyank Faldu, Univ. of Edinburgh. PhD thesis: “Addressing Variability in Reuse Prediction for Last-Level Caches,” 2014. First employment: ARM
 - Stavros Volos, EPFL. PhD thesis: “Memory Systems and Interconnects for Scale-Out Servers,” 2015. First employment: Microsoft Research.
 - Cansu Kaynak, EPFL. PhD thesis: “Shared Frontend for Manycore Server Processors,” 2015. First employment: Oracle Labs.
 - Onur Kocberber, EPFL. PhD thesis: “Accelerators for Data Processing,” 2015. First employment: Oracle Labs.
 - Sotiria Fytraki, EPFL. PhD thesis: “Architectural Support for Fine-Grained Program Monitoring,” 2014. First employment: Maxeler Technologies.
 - Pejman Lotfi-Kamran, EPFL. PhD thesis: “Scale-Out Processors,” 2013.
- Post-doctoral researchers

- Rakesh Kumar, Univ. of Edinburgh, Jan 2015 – Jun 2017 (co-supervised with Dr. Vijay Nagarajan). First employment: Assistant Professor at Norwegian University of Science and Technology (NTNU).
- Cheng-Chieh Huang, Univ. of Edinburgh, Jan 2015 – May 2017 (co-supervised with Dr. Vijay Nagarajan). First employment: TSCM