Date: January 22, 2024

RESUME

PERSONAL DETAILS

Full name:	Mark Silberstein
Phone:	+972-54-2106761
Email:	mark@ee.technion.ac.il
Web site:	https://marksilberstein.com

ACADEMIC DEGREES

2004-2010 **Ph.D. in Computer Science**
Thesis: A Distributed System for Genetic Linkage Analysis
Advisors: Prof. Dan Geiger, Prof. Assaf Schuster
Department of Computer Science, Technion, Haifa, Israel

1997-2001 **B.Sc. in Computer Engineering**, *with distinction*
Department of Electrical Engineering, Technion, Haifa, Israel

ACADEMIC APPOINTMENTS

2019-now **Associate Professor with tenure**
Department of Electrical and Computer Engineering
Department of Computer Science (Secondary appointment)
Technion, Haifa,Israel

2016 **A*STAR Visiting Professor**
Data Storage Institute, A-STAR, Singapore

2013-2019 **Assistant Professor**
Department of Electrical Engineering, Technion, Haifa, Israel

2011-2013 **Post doctoral fellow**
Host: Prof. Emmett Witchel, Operating Systems and Architecture group
Department of Computer Science, University of Texas at Austin, USA

2010-2011 **Post doctoral fellow**
Host: Prof. Idit Keidar
Department of Electrical Engineering, Technion, Haifa, Israel

2009 **Visiting scholar**
Host: Prof. Satoshi Matzuoka, Supercomputing group
Department of Computer Science, Tokyo Institute of Technology, Japan

2007 **Visiting scholar**
Host: Prof. John D. Owens, Visual Computing group
Department of Electrical and Computer Engineering
University of California at Davis, USA

PROFESSIONAL EXPERIENCE (OUTSIDE ACADEMIA)

2023-now **Consultant** *Pliops*

2021-2022 **Consultant** *Intel*

2021-2022 **Consultant** *NeuReality*

2018-2019 **Consultant** *Team-8*

2017-2018 **Consultant** *IBM Haifa Research Labs*

2011 **Research intern** *Microsoft Research, Redmond*

2000-2005 **Research intern** *IBM Haifa Research Labs, Distributed Systems Group*

2001-2003 **Military service** *High Performance and Cluster Computing group*
Armament division, Israel Defense Forces

1998-2000 **Software development intern** *Qualcomm Israel*

RESEARCH INTERESTS

My research spans a broad range of topics in computer systems, focusing on operating systems, programming models, applications, and hardware design for *hardware-accelerated computer systems*. My work lies in the intersection of Operating Systems, Computer Architecture, Systems Security, and Computer Networks.

TEACHING EXPERIENCE

Lecturer

Spring 2017-now [new, undergraduate+graduate] 046278: Accelerators and accelerated systems.

Winter 2015-now [undergraduate+graduate] 046209: Operating Systems Architecture

Winter 2020 - now [new, graduate] 048080: Advanced Topics: Operating systems and Secure Hardware

Winter 2021 [undergraduate, revised] 044101: Introduction to Computer Systems.

Spring 2020 [new, undergraduate+graduate] 046280: Principles of Cybersecurity (together with Prof. Eyal).

Spring 2018 [new, graduate] 048885: Advanced topics: Hardware and Software in post-CMOS era. New course developed and taught jointly with two other faculty.

Winter 2014-2019 [new, graduate] 048961: Advanced topics: Operating Systems Design and Implementation

Spring 2014-2016 [new, undergraduate+graduate] 046274: Advanced topics: GPU-accelerated systems

Spring 2015 [new, graduate] 048661: Advanced topics: Design and Implementation of Deep Learning Systems

Winter 2011 [undergraduate+graduate] 236370: Concurrent and Distributed Programming

Spring 2010 [undergraduate+graduate] 236308: Seminar in computer science: GPU computing

TECHNION ACTIVITIES

2023-now Advisory Committee for Cyber Security

2023 Special Investigation Committee of the Feb'23 Cyber Attack

2023-now Feb'23 Cyber Attack Recovery Committee

2020-now Technion Computing Infrastructure Committee

2018 Co-chair and organizer: The Summer School on Cyber Security: Hardware security and side channels. About *150* daily participants over three days.

2018-2021 Technion Cyber Security Research Center, Head of the Scientific Committee

DEPARTMENTAL ACTIVITIES

2020-now Head of the IT committee

2018-2021 Undergraduate studies committee

2016-2020 IT committee

2014-now Head of the Joint Computer Engineering Track committee

2014-now Undergraduate students consultant for the Computer Engineering Track

PUBLIC PROFESSIONAL ACTIVITIES

2023 Israeli Science Foundation, reviewer

2023 European Research Council, reviewer

2018 ACM SIGOPS selection committee for the Needham award for the Best PhD thesis in Computer Systems in Europe.

2017-2019 Official Blog of ACM SIG on Computer Architecture (SIGARCH), invited to maintain the column on research trends in operating systems.

2017-2029 *co-editor*, ACM Operating System Review Journal

2011-2017 Reviewer for scientific journals: IEEE Transactions on Parallel and Distributed Systems, IEEE Micro, Future Generation Computer Systems, IEEE Internet Computing, Journal of Artificial Intelligence Research, Concurrency and Computation - Practice and Experience, International Journal of Approximate Reasoning, IEEE Transactions on Computers, ACM Transactions on Information and System Security, ACM Transactions on Computer Systems

2013 National Science Foundation Review panel for CNS grants

2011 Israel Prime Minister Office Committee member on High Performance Infrastructures in Israel

PARTICIPATION IN ORGANIZING CONFERENCES

Conference Program Committees

2024 ASPLOS'25, USENIX ATC'25

2024 EuroSys'24 – *PC co-chair*, SOSP'24

2023 ASPLOS'23, EuroSys'23, OSDI'23, ASPLOS'23 – *Workshop chair*

2022 ASPLOS'22, USENIX ATC'22, SysTEX'22, SPMA'22, MICRO'22

2021 USENIX ATC'21, EuroSys'21, SOSP'21

2020 ASPLOS'20, VEE'20, SPMA Workshop (co-located with EuroSys'20) – *PC chair and organizer*

2019 SOSP'19, ATC'19, VEE'19, SFMA Workshop (co-located with EuroSys'19) – *co-organizer and PC chair*, European Doctoral Workshop (co-located with EuroSys'19)– *PC chair*, ASPLOS'19, PPoPP'19

2018 ASPLOS'18, SFMA'18 (co-located with EuroSys'18) – *co-organizer and PC chair*, Workshop on Advanced Memory Systems (co-located with ASPLOS'18) – *co-organizer*, SysTEX'18 Workshop (co-located with CCS'18) – *PC co-chair*, European Doctoral Workshop (co-located with EuroSys'18)

2017 SysTEX'17 workshop (co-located with SOSP'17), SOSP'17, Middleware, MaRS workshop (co-located with EuroSys'17) – *co-organizer and PC chair*

2016 SYSTOR'16 – *PC chair*, VEE'16, PACT'16, MaRS'16 workshop (co-located with EuroSys'16)

2015 SYSTOR'15, SFMA workshop (co-located with EuroSys'15), EuroSys'15

2014 ASPLOS'14, SYSTOR'14

2013 SOSP'13 – Poster Session, SYSTOR'13, CCGRID'13, IPDPS'13

MEMBERSHIP IN PROFESSIONAL SOCIETIES

USENIX, IEEE, ACM

FELLOWSHIPS, AWARDS AND HONORS

2023 Morton and Beverley Rechler Prize for Excellence in Research

2023 IEEE Micro Top Picks 2022

2023 IEEE Micro Top Picks 2022 Honorable Mention

2021 IBM Global University Program Academic Award

2020 First Prize Winner, Cyber Security Awareness Worldwide (CSAW) Regional Competition (Europe), best applied security paper

2020 Third Prize Winner, Cyber Security Awareness Worldwide (CSAW) Regional Competition (Israel), best applied security paper

2020 EuroSys Jochen Liedtke Young Researcher Award

2019 IEEE Micro Top Picks 2019

2018 First Prize Winner, Cyber Security Awareness Worldwide (CSAW) Regional Competition, best applied security paper

2018 Third Prize Winner, Cyber Security Awareness Worldwide (CSAW) Regional Competition, best applied security paper

2016 Best paper award, 6th International Workshop on Runtime and Operating Systems for Supercomputers

2014 Best paper award, 7th International Conference on Systems and Storage (SYSTOR)

2013 Horev Innovation Fellow

2013 Yahoo! ACE award

2013 Best Paper Award Runner-up, ASPLOS'13

2011 Best paper award, 4th International Conference on Systems and Storage (SYSTOR)

2010 Viterbi post doctoral fellowship

2010 Second Prize Winner, IEEE International Scalable Computing Challenge, IEEE Computer Society Technical Committee on Scalable Computing

2009 George Michael Memorial HPC Fellowship 2009 Honorable Mention Award, International Conference on High Performance Computing, Networking, Storage and Analysis, 2009

2008 Gutwirth graduate fellowship, Technion

2007 Jacobs-Qualcomm graduate fellowship, Technion

2006 Fein graduate fellowship, Faculty excellence award fellowship, Technion

2001 Alfaworks development award, IBM

GRADUATE STUDENTS

Completed PhD theses (Supervisor))

2015 Sangman Kim, "Operating System Support for Server Applications with Parallelism", with Prof. Emmett Witchel, UT Austin, (now in Apple)

2021 Oleksii Oleksenko, "Securing Trusted Execution Environments against Side Channel Attacks", with Prof. Christof Fetzer, TU Dresden, (now in Microsoft Research, Cambridge)

2021 Meni Orenbach, "Operating System Support for Trusted Execution Environments", (now in NVIDIA research)

2022 Haggai Eran, "Offloading computations to next-generation networking devices", (now in NVIDIA)

2022 Shai Bergman, "Operating system support for memory disaggregation", (now in Huawei Research Zurich)

Completed MSc theses (Supervisor)

2016 Sagi Shahar, "Efficient I/O operations on GPGPU devices", (now in Google)

2016 Feras Daoud, "High performance low latency networking from GPUs", (now in Mellanox)

2017 Amir Wated, "High-concurrency servers on GPUs", (now in Cisco)

2017 Shai Bergman, "High performance disk I/O on GPUs", (now PhD)

2018 Vasileos Dimitsas, "I/O prefetching for GPUs"

2018 Marina Minkin, "Securing shielded execution environments against side channel attacks", (now PhD at University of Michigan)

2019 Tanya Brokman, "Unified Heterogeneous Weakly-consistent Page Cache", (now in Microsoft)

2019 Lior Zeno, "Efficient I/O with accelerators", (now PhD)

2019 Maroun Tork, "Accelerator-centric architecture with SmartNICs", (now in Facebook)

2021 Menachem Edelman, "Accelerating neural network training via tensor sampling", (now in Intel)

2021 Lina Maudlej, "Operating system abstractions for accelerator disaggregation" (no in SpeedData)

2022 Alon Berkenstadt, "Automatic detection and mitigation of Spectre V2 attacks"

2023 Igor De Paula, "Using neural nets in network switches and routers"

Completed MSc theses (Co-supervisor)

2011 Uri Verner, "Processing real time data streams on accelerated systems", with Assaf Schuster and Avi Mendelson, (now in General Motors)

2016 Iftah Peretz, "Accelerating natural language parsing on GPUs", with Roei Reinchart

2016 Oren Ierushalmi, "File system aware FPGA", with Yoav Etzion

2016 Matan Hamilis, "Parallel additive fast Fourier transform algorithms", with Eli Ben-Sasson

2018 Arie Schwartz, "Medical monitoring using multi-modal cameras", with Guy Gilboa

PhD theses in progress

2019-2024 Lior Zeno, "In-network computing on programmable switches"

2019-2024 Alon Rashelbach, "Neural networks for packet classification"

MSc theses in progress

2022-2024 Saji Khashab, "Programming networks via packet programs"

2022-2024 Ron Marcus, "Using neural nets to accelerate solid state drives"

2023-2025 Assaf Klein, "Automatic detection of speculating execution leaks in CPU design"

2023-2025 Ori Ben Tzur, "Optimizing nested virtualization with Translation Pass Through"

RESEARCH GRANTS

Competitive

2022 ISF Using neural networks to accelerate network, storage and more. ($90K annual, 4 years)

2019 Prime Minister Bureau for Cyber Security Accelerated processing of data streams: Co-PI with A. Schuster ($35K annual, 3 years)

2018 ISF Operating systems for omni-programmable architectures, ($75K annual, 4 years)

2018 Prime Minister Bureau for Cyber Security Using power side channels for attacks and defenses, ($45K annual, 3 years)

2018 Prime Minister Bureau for Cyber Security: Co-PI with E. Biham Trusted execution environments: security and performance, ($15K annual, 3 years)

2016 ISF/China: Co-PI with S. Frankel Large-scale parallel computations on heterogeneous many-core supercomputers for LES of aerodynamics, ($30K annual, 3 years)

2014 ISF Operating system services for highly-threaded processors, ($75K annual, 4 years)

2014 NSF, UT Austin subcontract Harnessing highly-threaded accelerators for server workloads, ($15K annual, 4 years)

2013 MOST Advanced computing and cyber security, Co-PI with Y. Cassuto and R. Ginosar Secure compute-storage architectures, ($35K annual, 3 years)

Industrial

2022 KAMIN New hardware architecture for Longest Prefix Match ($130K, 2 years)

2023 Intel Gift: Secure acceleration on SmartNICs ($140K, 1 year)

2021 Intel Gift ($48K)

2021 IBM Academic Award ($40K)

2021 Western Digital Using novel storage devices in data centers ($55K)

2021 Intel Equipment grant ($20K)

2020 NVIDIA Equipment contribution ($40K)

2020 Huawei Accelerating Asynchronous Distributed Training of DNNs: Co-PI with A. Schuster, ($100K annual, 2 years)

2018 MAGNETon An infrastructure for application acceleration on smart network devices, ($120K annual, 2 years with Mellanox)

2018 MAGNET New architectures for multi-accelerator management, ($65K annual, 3 years)

2018 Huawei OS support for omni-programmable systems, ($75K annual, 1 year)

2017 Intel Compiler support for efficient execution of unmodified applications in SGX enclaves, ($50K annual, 1 year)

2015 ICRI-CI Distributed accelerated system for deep learning using Intel Xeon-Phi, ($25K annual, 2 years)

2014 MAGNET High Performance VLSI technologies, Operating system support for GPUs, ($68K annual, five years)

PLENARY, KEYNOTE AND INVITED TALKS (SINCE 2014)

2014 **Operating system services for high throughput processors**,

(invited) Computer Engineering Club, Technion; (invited) Yahoo! Research, Haifa, Israel; (invited) PRACE - Partnership for Advanced Computing In Europe, Tel Aviv, Israel; (invited) Max-Plank Institute of Software Systems, Kaiserslautern, Germany

2015 **Operating system services for high throughput processors**,

(invited) Computer Science Colloquium, Cornell University, USA; (invited) IBM Watson Research Center, Yorktown Heights, USA; (plenary) DevelopEx Conference, Israel; (invited) Memorial Merlin Lectures, Technion

2015 **Accelerator-centric operating system: rethinking the role of CPUs in modern computer systems**, (plenary) Technion Computer Engineering Conference, Israel

2015 **Series of lectures on GPU computing**, three-day summer school, A*Star Data Storage Institute, Singapore

2016 **Never trust your GPU**, (plenary) Technion Cyber Security Inauguration Workshop; (invited) Tel Aviv University, Security Colloquium;

2016 **Accelerator-centric operating system: rethinking the role of CPUs in modern computer systems**, (plenary and panel) International Workshop on Runtime and Operating Systems for Supercomputers (ROSS)

2016 **ActivePointers: a case for software address translation on GPUs**, (invited) University de Catalunia, Barcelona, Spain

2016 **Operating system services for high throughput processors**, (invited) University of Californa, Davis, USA

2016 **FPGAs: A game changer for machine learning workloads or Nothing new under the Sun** (plenary), ICRI-CI Retreat and Conference, Intel, Israel

2016 **Accelerator-centric operating systems architecture**, (invited) NSF sponsored Workshop on Architecture and Software for Emerging Applications (WASEA), co-located with PACT, Haifa, Israel

2016 **Parallel stochastic gradient descent: the case for native GPU-side GPI**, InterTWINE EU project meeting, Manchester, UK

2016 **GPU-to-GPU networking: past, present, future**, High (invited) Performance Computing Workshop, Mechanical Engineering Department, Technion

2017 **Accelerated future of computer systems**, Israel Challenge, Technion

2017 **Never trust your GPU**, (invited) Ben Gurion University, Cyber Security seminar

2017 **OmniX: Accelerator-centric OS design**, (invited) International Workshop "Beyond CMOS: from devices to systems"

2017 **The power of software address translation with applications to SGX**, (**keynote**), International Workshop on Secure Execution Environments (SYSTEX - co-located with SOSP17, Shanghai, China)

2017 **GAON: general application offload to near-network processors**, Cambridge University, Microsoft Research Cambridge, Imperial College London, University of Wisconsin Madison, University of Michigan Ann Arbor, Georgia Tech, University of Washington, Seattle.

2018 **Adaptable security**, (invited) Israel-France Cyber-security Forum

2018 **What if your phone's battery could talk**, (invited) Cyber-security Summer School, Forum on security and tools against terror.

2018 **Foreshadow: speculative attacks on Intel SGX**, Cornell-Tech New York, UT Austin

2018 **NICA: Accelerating applications on SmartNICs**, Cornell, Ithaca

2019 **When SGX fell victim to speculative execution attacks**, (invited) Technion Cyber-day, Technion; TU Dresden

2019 **OmniX: Accelerator-centric OS design**, Huawei International Technical Workshop (invited), TU Dresden, EPFL Switzerland (invited)

2019 **NICA: Accelerating applications on SmartNICs**, Hebrew University Jerusalem

2021 **Omnix: Accelerator-centric OS design**, Pliops (invited)

2021 **Computational approach to packet classification**, Intel, NVIDIA, Rice University, Hebrew University, Intel Labs, Intel Barefoot Networks, IBM Haifa Research Labs

2022 **Securing Trusted Execution Environments from side-channel and untrusted interface attacks**, Microsoft Research Cambridge

2022 **Cloud networks: trends, challenges, solutions**, TU Darmstadt

2022 **Computational approach to packet classification**, Global Networking forum, Tel Aviv

2022 **Rethinking OS optimizations for memory disaggregation**, Storage forum, Tel Aviv

2023 **Computational approach to packet classification**, KTH, University of British Columbia

2023 **Multi-tenant in-network computing**, NVIDIA

2024 **Securing Trusted Execution Environments from side-channel and untrusted interface attacks**, Tel Aviv University

2024 **Automatic detection of speculative execution attacks**, Hebrew University

PUBLICATIONS

Student authors under my supervision are marked by asterisk

Theses

T-1 M. Silberstein. *A Distributed System for Genetic Linkage Analysis*. PhD Thesis, Technion, 2010

Referred papers in professional journals

J-1 M. Silberstein, A. Tzemach, N. Dovgolevsky, M. Fishelson, A. Schuster, and D. Geiger. On-line system for faster linkage analysis via parallel execution on thousands of personal computers. *Americal Journal of Human Genetics*, 78(6):922–935, 2006

J-2 M. Silberstein, O. Weissbrod, L. Otten, A. Tzemach, A. Anisenia, O. Shtark, D. Tuberg, E. Galfrin, I. Gannon, and A. Shalata. A system for exact and approximate genetic linkage analysis of snp data in large pedigrees. *Bioinformatics*, 29(2):197–205, 2013

J-3 M. Silberstein. GPUs: high performance accelerators for parallel applications. *Ubiquity Symposium: The Multicore Transformation*, pages 1:1–1:13, 2014

J-4 M. Silberstein, B. Ford, and E. Witchel. A Case for Operating System Abstractions on GPUs. *Communications of ACM*, 57/No.12:68–79, 2014

J-5 M. Silberstein, B. Ford, I. Keidar, and E. Witchel. GPUfs: Integrating a File System with GPUs. *ACM Transactions on Computer Systems (TOCS)*, 32(1):1:1–1:31, Feb. 2014

J-6 M. Silberstein, S. Kim*, S. Huh, X. Zhang, Y. Hu, A. Wated*, and E. Witchel. GPUnet: Networking Abstractions for GPU Programs. *ACM Transactions on Computer Systems (TOCS)*, 34(3):9:1–9:31, Sept. 2016

J-7 S. Shahar*, S. Bergman*, and M. Silberstein. ActivePointers: A Case for Software Address Translation on GPUs. *SIGOPS Operating Systems Review*, 52(1):84–95, Aug. 2018

J-8 J. V. Bulck, M. Minkin*, O. Weisse, D. Genkin, B. Kasikci, F. Piessens, M. Silberstein, T. F. Wenisch, Y. Yarom, and R. Strackx. Breaking Virtual Memory Protection and the SGX Ecosystem with Foreshadow. *IEEE Micro*, 39(3):66–74, 2019

J-9 S. Bergman*, T. Brokhman*, T. Cohen, and M. Silberstein. SPIN: Seamless Operating System Integration of Peer-to-Peer DMA Between SSDs and GPUs. *ACM Trans. Comput. Syst.*, 36(2):5:1–5:26, Apr. 2019

J-10 A. Rashelbach*, O. Rottenstreich, and M. Silberstein. A computational approach to packet classification. *IEEE/ACM Transactions on Networking*, 30(3):1073–1087, 2022

J-11 S. Bergman*, N. Cassel, M. Bjørling, and M. Silberstein. ZNSwap: Un-Block Your Swap. *ACM Transactions on Storage*, 19(2), 2023

J-12 A. Rashelbach*, O. Rottenstreich, and M. Silberstein. Scaling by learning: Accelerating open vswitch data path with neural networks. *IEEE/ACM Transactions on Networking*, 31(3):1230–1243, 2023

J-13 O. Oleksenko, B. Köpf, C. Fetzer, and M. Silberstein. Revizor: Testing Black-box CPUs against Speculation Contracts. *IEEE Micro Top Picks: selected papers from the 2022 Computer Architecture Conferences*, 2023

Book chapters

B-1 M. Silberstein, A. Schuster, and J. Owens. Applying Software Managed Caching and CPU-GPU Scheduling for Accelerating Dynamic Computations . In W. mei W. Hwu, editor, *GPU Computing Gems Jade Edition*, pages 501 – 519. Morgan Kaufmann, 2011

Refereed papers in conference proceedings

C-1 M. Silberstein, D. Geiger, A. Schuster, and M. Livny. Scheduling Mixed Workloads in Multi-grids: The Grid Execution Hierarchy. In *Proc. of the International Symposium on High Performance Distributed Computing*, HPDC'15, pages 291–302. ACM/IEEE, 2006

C-2 V.Kravtsov, D.Carmeli, A.Schuster, B.Yoshpa, M.Silberstein, W.Dubitzky. Quasi-Opportunistic Supercomputing in Grid Environments. In *International Conference on Algorithms and Architectures*, ICAA, pages 233 – 244. IEEE, 2008

C-3 M. Silberstein, A. Schuster, D. Geiger, A. Patney, and J. Owens. Efficient Computation of Sum-Products on GPUs Through Software-Managed Cache. In *International Conference on Supercomputing*, ICS'08, pages 309–318. ACM, 2008

C-4 M. Silberstein, A. Sharov, D. Geiger, and A. Schuster. GridBot: Execution of Bags of Tasks in Multiple Grids, (George Michael Memorial HPC Fellowship 2009 Honorable Mention Award). In *International Conference for High Performance Computing, Networking, Storage and Analysis*, SC'09, pages 11:1–11:12. ACM, 2009

C-5 M. Silberstein. Building an Online Domain-Specific Computing Service over Non-dedicated Grid and Cloud Resources: The Superlink-Online Experience. In *International Symposium on Cluster, Cloud and Grid Computing*, CCGRID'11, pages 174–183. IEEE, 2011

C-6 U. Verner*, A. Schuster, and M. Silberstein. Processing Data Streams With Hard Real-time Constraints on Heterogeneous Systems. In *International Conference on Supercomputing*, ICS'11, pages 120–129. ACM, 2011

C-7 M. Silberstein and N. Maruyama. An Exact Algorithm for Energy-Efficient Acceleration of Task Trees on CPU/GPU Architectures, (Best Paper Award). In *International Conference on Systems and Storage*, SYSTOR'04, pages 7:1–7:7. ACM, 2011

C-8 C. J. Rossbach, J. Currey, M. Silberstein, B. Ray, and E. Witchel. PTask: Operating System Abstractions To Manage GPUs as Compute Devices . In *Symposium on Operating Systems Principles*, SOSP'22, pages 233–248. ACM, 2011

C-9 U. Verner*, A. Schuster, M. Silberstein, and A. Mendelson. Scheduling of Real-Time Data Streams on Heterogeneous Multi-GPU Systems. In *International Systems and Storage Conference*, SYSTOR'05, pages 7:1–7:12. ACM, 2012

C-10 O. Ben-Yehuda, M. Silberstein, A. Sharov, A. Iosup, and A. Schuster. ExPERT: Pareto-Efficient Task Replication on Grids and a Cloud. In *International Parallel and Distributed Processing Symposium*, IPDPS'12, pages 167 – 178. IEEE, 2012

C-11 A. M. Dunn, M. Z. Lee, S. Jana, S. Kim, M. Silberstein, Y. Xu, V. Shmatikov, and E. Witchel. Eternal Sunshine of the Spotless Machine: Protecting Privacy with Ephemeral Channels. In *Symposium on Operating Systems Design and Implementation*, OSDI'12, pages 61–75. USENIX, 2012

C-12 M. Silberstein, B. Ford, I. Keidar, and E. Witchel. GPUfs: Integrating File Systems with GPUs, *best paper runner-up*. In *International Conference on Architectural Support for Programming Languages and Operating Systems*, APSLOS'13, pages 485–498. ACM, 2013

C-13 M. Silberstein, L. Ganesh, Y. Wang, L. Alvisi, and M. Dahlin. Lazy Means Smart: Reducing Repair Bandwidth Costs in Erasure-coded Distributed Storage, *best paper award*. In *International Conference on Systems and Storage*, SYSTOR'7, pages 15:1–15:7. ACM, 2014

C-14 S. Kim*, S. Huh, X. Zhang, Y. Hu, A. Wated*, E. Witchel, and M. Silberstein. GPUnet: Networking Abstractions for GPU Programs. In *Symposium on Operating Systems Design and Implementation*, OSDI'14, pages 201–216. USENIX, 2014

C-15 A. Newell, G. Kliot, A. Gopalan, I. Menache, S. Akiyama, and M. Silberstein. Optimizing Distributed Actor Systems for Dynamic Interactive Services. In *European Conference on Computer Systems*, EuroSys'16, pages 38:1–38:15. ACM, 2016

C-16 S. Shahar* and M. Silberstein. Supporting Data-driven I/O on GPUs Using GPUfs. In *International Conference on Systems and Storage*, SYSTOR'9, pages 12:1–12:11. ACM, 2016

C-17 S. Shahar*, S. Bergman*, M. Silberstein. ActivePointers: the Case for Software Address Translation on GPUs. In *International Symposium on Computer Architectures*, ISCA'16, pages 596–608. ACM, 2016

C-18 M. Hamilis*, E. Ben-Sason, E. Tromer, and M. Silberstein. Accelerating Binary Finite Fields Multiplication on GPUs via Register Cache. In *ACM International Conference on Supercomputing* , ICS'16, pages 35:1–35:12. ACM, 2016

C-19 E. Ben-Sasson, I. Bentov, A. Chiesa, A. Gabizon, D. Genkin, M. Hamilis*, E. Pergament, M. Riabzev, M. Silberstein, E. Tromer, and M. Virza. Computational Integrity with a Public Random String from Quasi-Linear PCPs. In J.-S. Coron and J. B. Nielsen, editors, *Advances in Cryptology – EUROCRYPT 2017*, pages 551–579, Cham, 2017. Springer International Publishing

C-20 M. Orenbach*, P. Lifshits*, M. Minkin*, and M. Silberstein. Eleos: Exit-less Operating System Services for SGX Enclaves. In *European Conference on Computer Systems*, EuroSys'17, pages 238–253. ACM, 2017

C-21 M. Silberstein. Omnix: an Operating System for Omni-programmable Computer Systems. In *ACM Workshop on Hot Topics in Operating Systems*, HotOS, pages 69–75. ACM, 2017

C-22 S. Bergman*, T. Brokhman*, T. Cohen, and M. Silberstein. SPIN: Seamless Operating System Integration of Peer-to-Peer DMA Between SSDs and GPUs. In *2017 USENIX Annual Technical Conference (USENIX ATC 17)*, USENIX ATC'17, pages 167–179. USENIX, 2017

C-23 O. Oleksenko, B. Trach, R. Krahn, M. Silberstein, and C. Fetzer. Varys: Protecting SGX Enclaves from Practical Side-Channel Attacks. In *2018 USENIX Annual Technical Conference (USENIX ATC 18)*, USENIX ATC'18, pages 227–240. USENIX, 2018

C-24 P. Lifshits*, R. Forte*, Y. Hoshen, M. Halpern, M. Philipose, M. Tiwari, and M. Silberstein. Power to Peep-all: Inference Attacks by Malicious Batteries on Mobile Devices. In *Privacy Enhancing Technologies Symposium (PETS)*, volume 2018, pages 141 – 158. De Gruyter, 2018

C-25 J. V. Bulck, M. Minkin*, O. Weisse, D. Genkin, B. Kasikci, F. Piessens, M. Silberstein, T. F. Wenisch, Y. Yarom, and R. Strackx. Foreshadow: Extracting the Keys to the Intel SGX Kingdom with Transient Out-of-Order Execution, *Selected as Micro Top Picks '18*. In *USENIX Security Symposium*, USENIX Security'18, page 991–1008. USENIX, 2018

C-26 H. Eran*, L. Zeno*, I. Zsolt, and M. Silberstein. Design Patterns for Code Reuse in HLS Packet Processing Pipelines. In *IEEE International Symposium on Field-Programmable Custom Computing Machines*, FCCM'19. IEEE, 2019

C-27 H. Eran*, L. Zeno*, M. Tork*, G. Malka, and M. Silberstein. NICA: An Infrastructure for Inline Acceleration of Network Applications. In *2019 USENIX Annual Technical Conference (USENIX ATC 19)*, pages 345–362, Renton, WA, July 2019. USENIX Association

C-28 T. Brokman*, P. Lifshits*, and M. Silberstein. GAIA: Unified Page Cache for Heterogeneous Systems. In *2019 USENIX Annual Technical Conference (USENIX ATC 19)*, pages 661–674, Renton, WA, 2019

C-29 M. Orenbach*, Y. Michalevsky, C. Fetzer, and M. Silberstein. COSMIX: a Compiler Approach to Supporting Secure Page Faults in Trusted Execution Environments. In *2019 USENIX Annual Technical Conference (USENIX ATC 19)*, pages 555–570, Renton, WA, 2019

C-30 A. Watad*, A. Libov, O. Shaham, E. Bortinkov, and M. Silberstein. Achieving Scalability in a k-NN Multi-GPU Network Service with Centaur. In *2019 28th International Conference on Parallel Architectures and Compilation Techniques (PACT)*, pages 245–257, 2019

C-31 O. Oleksenko*, B. Trach, M. Silberstein, and C. Fetzer. SpecFuzz: Bringing Spectre-type Vulnerabilities to the Surface. In *29th USENIX Security Symposium (USENIX Security 20)*, pages 1481–1498. USENIX Association, Aug. 2020

C-32 M. Tork*, L. Maudlej*, and M. Silberstein. Lynx: a SmartNIC-driven Accelerator-centric Architecture for Network Servers. In *The 25th International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS '20, page 117–131. ACM, 2020

C-33 M. Orenbach*, A. Baumann, and M. Silberstein. Autarky: Closing Controlled Channels with Self-paging Enclaves. In *The 15th European Conference on Computer Systems*, EuroSys '20. ACM, 2020

C-34 A. Rashelbach*, O. Rottenstreich, and M. Silberstein. A Computational Approach to Packet Classification. In *Annual Conference of the ACM Special Interest Group on Data Communication on the Applications, Technologies, Architectures, and Protocols for Computer Communication*, SIGCOMM '20, page 542–556, 2020

C-35 S. Eliad, I. Hakimi, A. D. Jagger, M. Silberstein, and A. Schuster. Fine-tuning Giant Neural Networks on Commodity Hardware with Automatic Pipeline Model Parallelism. In *2021 USENIX Annual Technical Conference (USENIX ATC 21)*, pages 381–396. USENIX Association, July 2021

C-36 M. Adelman*, I. Hakimi, K. Levi, and M. Silberstein. Faster Neural Network Training via Approximate Tensor Operations. In *2021 Conference on Neural Information Processing Systems (NeurIPS'21)*, page 1–8, 2021

C-37 O. Oleksenko*, C. Fetzer, B. Köpf, and M. Silberstein. Revizor: Testing Black-box CPUs against Speculation Contracts, *Selected as Micro Top Picks '22*. In *International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS '22, page 1–16. ACM/IEEE, 2022

C-38 H. Eran*, M. Fudim, G. Malka, G. Shalom, N. Cohen, A. Hermony, D. Levi, L. Liss, and M. Silberstein. FlexDriver: A Network Driver for Your Accelerator, *Selected as Micro Top Picks '22 honorable mention*. In *International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS '22, page 1–16. ACM/IEEE, 2022

C-39 A. Rashelbach*, O. Rottenstreich, and M. Silberstein. Scaling Open vSwitch with a Computational Cache. In *2022 USENIX Network Systems Design and Implementation*, NSDI '22, page 1–18. USENIX Association, 2022

C-40 C. Raiciu, V. Olteanu, A. Popa, M. Handley, H. Eran*, M. Silberstein, D. Dumitrescu, C. Baciu, and G. Nikolaidis. An Edge-queued Datagram Service for All Data-center Traffic. In *2022 USENIX Network Systems Design and Implementation*, NSDI '22, page 1–16. USENIX Association, 2022

C-41 L. Zeno*, D. Ports, J. Nelson, D. Kim, S. L. Feibish, I. Keidar, A. Rinberg, A. Rashelbach*, I. De-Paula*, and M. Silberstein. SwiSh: Distributed Shared State Abstractions for Programmable Switches. In *2022 USENIX Network Systems Design and Implementation*, NSDI '22, page 1–15. USENIX Association, 2022

C-42 L. Vilanova, L. Maudlej*, S. Bergman*, T. Miemietz, M. Hille, N. Asmussen, M. Roitzsch, H. Härtig, and M. Silberstein. Slashing the Disaggregation Tax in Heterogeneous Data Centers with FractOS. In *The 17th European Conference on Computer Systems*, EuroSys '22, page 352–367, New York, NY, USA, 2022. ACM

C-43 S. Bergman*, P. Faldu, B. Grot, L. Vilanova, and M. Silberstein. Reconsidering OS Memory Optimizations in the Presence of Disaggregated Memory. In *The 2022 ACM SIGPLAN International Symposium on Memory Management*, ISMM '22, page 1–14, New York, NY, USA, 2022. Association for Computing Machinery

C-44 S. Bergman*, N. Cassel, M. Bjørling, and M. Silberstein. ZNSwap: un-Block your Swap. In *2022 USENIX Annual Technical Conference (USENIX ATC 22)*, pages 1–18, Carlsbad, CA, July 2022. USENIX Association

C-45 O. Oleksenko, B. Köpf, M. Guarnieri, and M. Silberstein. Hide and Seek with Spectres: Efficient Discovery of Speculative Information Leaks with Random Testing. In *IEEE Symposium on Security and Privacy*, S&P '23, page 1–16. IEEE, 2023

C-46 S. Bergman*, M. Silberstein, T. Shinagawa, P. Pietzuch, and L. Vilanova. Translation Pass-Through for Near-Native Paging Performance in VMs. In *2023 USENIX Annual Technical Conference (USENIX ATC 23)*, pages 753–768, Boston, MA, July 2023. USENIX Association

C-47 S. Constable, J. V. Bulck, X. Cheng, Y. Xiao, C. Xing, I. Alexandrovich, T. Kim, F. Piessens, M. Vij, and M. Silberstein. AEX-Notify: Thwarting Precise Single-Stepping Attacks through Interrupt Awareness for Intel SGX Enclaves. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 4051–4068, Anaheim, CA, Aug. 2023. USENIX Association

C-48 A. Rashelbach*, I. De-Paula*, and M. Silberstein. NeuroLPM: Scaling Longest Prefix Match Hardware with Neural Networks. In *56th IEEE/ACM International Symposium on Microarchitecture*, MICRO '56, pages 1–15. IEEE/ACM, 2023

C-49 S. Khashab*, A. Rashelbach*, and M. Silberstein. Multitenant In-Network Acceleration with SwitchVM. In *2024 USENIX Network Systems Design and Implementation*, NSDI '24, pages 1–16. USENIX Association, 2023

Refereed papers in workshop proceedings

W-1 M. Silberstein, M. Factor, and D. Lorenz. DYNAMO - DirectorY, Net Archiver and MOver. In *Third International Workshop on Grid Computing, in conjunction with SC02*, GRID'02, pages 256–267. ACM, 2002

W-2 M. Silberstein, D. Geiger, and A. Schuster. A Distributed System for Genetic Linkage Analysis. In *International Workshop on Distributed, High-Performance and Grid Computing in Computational Biology*, GCCB'06, pages 110–123, 2007

W-3 M. Silberstein, G. Kliot, A. Sharov, A. Schuster, and M. Livny. Materializing Highly Available Grids *(short paper)*. *Proc. of the International Symposium on High-Performance Distributed Computing*, pages 321–323, 2006

W-4 V.Kravtsov, D.Carmeli, A.Schuster, B.Yoshpa, M.Silberstein, W. Dubitzky. Quasi-Opportunistic Supercomputing in Grids *(short paper)*. In *International Symposium on High Performance Distributed Computing*, HPDC'07. ACM, 2007

W-5 M. Silberstein. Building an Online Computing Service Over Volunteer Grid Resources, *invited paper*. In *Workshop on Desktop Grids and Volunteer Computing Systems,in conjunction with IPDPS'11*, PCGRID, pages 1909–1917. IEEE, 2011

W-6 M. Silberstein, S. Kim*, S. Huh, Y. Hu, X. Zhang, and E. Witchel. GPUnet: networking abstractions for GPUs. In *The 4th Workshop on Systems for Future Multicore Architectures*, SFMA'4, 2014

W-7 L. Otten, R. Dechter, M. Silberstein, and D. Geiger. Maximum Likelihood Haplotyping through Parallelized Search on a Grid of Computers. In *International Conference on Research in Computational Molecular Biology*, RECOMB'09, 2009

W-8 S. Shahar* and M. Silberstein. Supporting Input-dependent Data on GPUs. In *The 5th Workshop on Systems for Future Multicore Architectures*, SFMA'5, 2015

W-9 L. Zeno*, A. Mendelson, M. Silberstein. The Case for I/O Preemption on Discrete GPUs. In *International Workshop on GPU computing systems*, GPGPU '9, pages 63–71, 2016

W-10 F. Daoud*, A. Watad*, M. Silberstein. GPUrdma: GPU-side Library for High Performance Networking from GPU kernels, *best paper award*. In *ACM International Workshop on Runtime and Operating Systems for Supercomputers*, ROSS '16, pages 6:1–6:8, 2016

W-11 S. Kim*, Z. Zhu, Y. Rozhanski*, E. Witchel, M. Silberstein. Understanding the security of discrete GPUs. In *International Workshop on GPU computing systems*, GPGPU '10, pages 1–11, 2017

W-12 A. Watad* and M. Silberstein. GPUmore: Scalable Multi-GPU Dataset Centric Network Servers. In *Workshop on Multi-core and Rack-scale Systems*, MARS, 2017

W-13 H. Eran* and M. Silberstein. NICA: OS support for Near-data Network Application Accelerators. In *Workshop on Multi-core and Rack-scale Systems*, MARS, 2017

W-14 H. Eran*, L. Liss, D. Levi, and M. Silberstein. NFV acceleration: the Role of the NIC. In *International Workshop on Systems for Future Multicore and Heterogeneous Architectures (SFMA)*, 2018

W-15 L. Vilanova*, Y. Etsion, and M. Silberstein. One Interface To Rule Them All: A Hardware/Software Co-design For Disaggregated Computing. In *International Workshop on Systems for Future Multicore and Heterogeneous Architectures (SFMA)*, 2019

W-16 L. Zeno*, D. R. K. Ports, J. Nelson, and M. Silberstein. SwiShmem: Distributed Shared State Abstractions for Programmable Switches. In *The 19th ACM Workshop on Hot Topics in Networks*, HotNets '20, page 160–167. ACM, 2020

W-17 A. Rashelbach*, O. Rottenstreich, and M. Silberstein. Nucleotide String Indexing using Range Matching (poster). In *14th ACM Conference on Bioinformatics, Computational Biology, and Health Informatics 2023*, ACM-BCB '14, page 1–8. ACM, 2023