Michael **Wei** Full-Stack Systems Researcher.

+1 626-367-3677
michael@wei.systems
http://wei.systems
no2chem

about

I'm a full—stack systems researcher. I dive deep into all aspects of complex systems to understand what makes them tick, and what makes them break. My research and open source contributions span multiple disciplines, including computer architecture, distributed systems, operating systems, and programming languages.

Through the course of my work, I have not only developed cost-effective hardware to simplify distributed systems, but also devised innovative approaches to streamline traditionally opaque operating system interfaces. My aim is always to leverage interdisciplinary insights for improved efficiency and resilience, while actively sharing knowledge through accessible, comprehensive documentation and meaningful open-source contributions.

experience

2023–Now **Senior Research Scientist, AI Labs** — VMware Research
As a researcher at the VMware Research Group, a part of VMware AI Labs, I lead efforts in sustainable AI systems infrastructure, leveraging low power and edge accelerators on platforms such as the AGX Orin.

2020–2023 **Senior Research Scientist** — VMware Research
I lead the clocks project, which aims to democratize precision time and leverage it as a distributed primitive. I also serve as a technical advisor to the VMware blockchain product, especially in scaling the VMware Ethereum stack.

2017–2020 **Research Scientist** — VMware Research
I published award-winning papers on a broad range of topics including bridging the semantic gaps between kernel and userspace, interfaces for programmable hardware, RDMA and security.

2015–2017 **Postdoctoral Researcher** — VMware Research
I led the tech transfer of CorfuDB, a high throughput, strongly consistent research database as open-source software with over 50 contributors into NSX, VMware's market leading multi-billion dollar SDN product. CorfuDB now serves as the distributed control plane of NSX.

Winter 2015 **Research Intern** — VMware Research
I led the rebuild of CORFU, a research project I worked on at Microsoft Research, from scratch in Java as CorfuDB, adding object support, in-memory time-travel and new transaction resolution mechanisms.

2011–2013 **Consultant Researcher** — Microsoft Research
I implemented and evaluated ZooKeeper over CORFU, one of the first scale-out applications running on top of the CORFU distributed log.

Summer 2011 **Research Intern** — Microsoft Research
I built a hardware prototype of a CORFU Flash Unit using the BEE3 FPGA Platform.

2010–2017 **Graduate Researcher** — University of California, San Diego
At the Non-Volatile Systems Laboratory (NVSL), I led work on reliably erasing flash memory and reported findings to the military and other government organizations.

2008–2009 **Undergraduate Researcher** — University of California, Irvine
At the Cognitive Anteater Robotics Laboratory (CARL), I worked on an optimized version of the Izhikevich spiking neuron model on the IBM CELL.

2006–Now **Independent Contractor** Wei Enterprises
I build hardware and software for startups under tight deadlines. My code runs in consumer products all over the world.

education

2010–2017 **Ph.D., Computer Science** University of California, San Diego
Advisor Steven Swanson
Thesis CORFU: A Platform for Scalable Consistency

2010–2015 **C. Phil, Computer Science** University of California, San Diego

2010–2012 **M.S., Computer Science** University of California, San Diego

2004–2009 **B.S., Computer Science** University of California, Irvine

2004–2009 **B.S., Biological Sciences** University of California, Irvine

2004–2009 **B.A., Philosophy** University of California, Irvine

awards

2022 **Best Paper Award** NSDI 2022
Graham: Synchronizing Clocks by Leveraging Local Clock Properties

2020 **Best Paper Award** Eurosys 2020
Don't shoot down TLB shootdowns!

2019 **Best Paper Award** Systor 2019
Storm: A Fast Transactional Dataplane for Remote Data Structures

2017 **Best Paper Award** USENIX ATC 2018
The Design and Implementation of Hyperupcalls

2016 **Best Paper Award** USENIX ATC 2016
Replex: A Scalable, Highly Available Multi-Index Store

2014 **Best WACI Talk** ASPLOS 2014
Dirty RAM and Rotten Caches: Saving the World From Useless Updates

2012 **Google Security Hall of Fame** Google
Google Voice Authentication Flaw

2012 **Graduate Research Fellowship Program (GRFP)** NSF
Fellow Award Receipient

2011 **Graduate Research Fellowship Program (GRFP)** NSF
Honorable Mention

publications

Refereed Papers

NSDI '22 Graham: Synchronizing Clocks by Leveraging Local Clock Properties
Best Paper Ali Najafi and **Michael Wei**
19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22), 2022, Renton, WA

SYSTOR '21 Dealing with (Some of) the Fallout from Meltdown
Nadav Amit, **Michael Wei**, and Dan Tsafrir
Proceedings of the 14th ACM International Conference on Systems and Storage, 2021, Haifa, Israel. DOI: *10.1145/3456727.3463776*

HotOS '21 — Systems Research is Running out of Time
Ali Najafi, Amy Tai, and **Michael Wei**
Proceedings of the Workshop on Hot Topics in Operating Systems, 2021, Ann Arbor, Michigan. DOI: *10.1145/3458336.3465293*

USENIX ATC '21 — RainBlock: Faster Transaction Processing in Public Blockchains
Soujanya Ponnapalli, Aashaka Shah, Souvik Banerjee, Dahlia Malkhi, Amy Tai, Vijay Chidambaram, and **Michael Wei**
2021 USENIX Annual Technical Conference (USENIX ATC 21), 2021

OSDI '21 — Optimizing Storage Performance with Calibrated Interrupts
Amy Tai, Igor Smolyar, **Michael Wei**, and Dan Tsafrir
15th USENIX Symposium on Operating Systems Design and Implementation (OSDI 21), 2021

EuroSys '20 **Best Paper** — Don't Shoot down TLB Shootdowns!
Nadav Amit, Amy Tai, and **Michael Wei**
Proceedings of the Fifteenth European Conference on Computer Systems, 2020, Heraklion, Greece. DOI: *10.1145/3342195.3387518*

USENIX ATC '19 — JumpSwitches: Restoring the Performance of Indirect Branches In the Era of Spectre
Nadav Amit, Fred Jacobs, and **Michael Wei**
2019 USENIX Annual Technical Conference (USENIX ATC 19), 2019, Renton, WA

SYSTOR '19 **Best Paper** — Storm: A Fast Transactional Dataplane for Remote Data Structures
Stanko Novakovic, Yizhou Shan, Aasheesh Kolli, Michael Cui, Yiying Zhang, Haggai Eran, Boris Pismenny, Liran Liss, **Michael Wei**, Dan Tsafrir, and Marcos Aguilera
Proceedings of the 12th ACM International Conference on Systems and Storage, 2019, Haifa, Israel. DOI: *10.1145/3319647.3325827*

ASPLOS '19 — Just-In-Time Compilation for Verilog: A New Technique for Improving the FPGA Programming Experience
Eric Schkufza, **Michael Wei**, and Christopher J. Rossbach
Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 19), 2019, Providence, RI, USA. DOI: *10.1145/3297858.3304010*

USENIX ATC '18 — Remote regions: a simple abstraction for remote memory
Marcos K. Aguilera, Nadav Amit, Irina Calciu, Xavier Deguillard, Jayneel Gandhi, Stanko Novaković, Arun Ramanathan, Pratap Subrahmanyam, Lalith Suresh, Kiran Tati, Rajesh Venkatasubramanian, and **Michael Wei**
2018 USENIX Annual Technical Conference (USENIX ATC 18), 2018, Boston, MA

USENIX ATC '18 **Best Paper** — The Design and Implementation of Hyperupcalls
Nadav Amit and **Michael Wei**
2018 USENIX Annual Technical Conference (USENIX ATC 18), 2018, Boston, MA

OSDI '18 — Sharing, Protection, and Compatibility for Reconfigurable Fabric with AmorphOS
Ahmed Khawaja, Joshua Landgraf, Rohith Prakash, **Michael Wei**, Eric Schkufza, and Christopher J. Rossbach
13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18), 2018, Carlsbad, CA

HotOS '17 — Hypercallbacks: Decoupling Policy Decisions and Execution
Nadav Amit, **Michael Wei**, and Cheng-Chun Tu
Proceedings of the 16th Workshop on Hot Topics in Operating Systems, 2017, Whistler, BC, Canada. DOI: *10.1145/3102980.3102987*

SysTeX '17 — Hypercallbacks: A New Mechanism for Trusted, Secure Introspection
Michael Wei and Nadav Amit
2nd Workshop on System Software for Trusted Execution (SysTEX 2017), 2017, Shanghai, CN

NSDI '17 — vCorfu: A Cloud-Scale Object Store on a Shared Log
Michael Wei, Amy Tai, Christopher J. Rossbach, Ittai Abraham, Maithem Munshed, Medhavi Dhawan, Jim Stabile, Udi Wieder, Scott Fritchie, Steven Swanson, Michael J. Freedman, and Dahlia Malkhi
14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17), 2017, Boston, MA

USENIX ATC '16 **Best Paper**
Replex: A Scalable, Highly Available Multi-Index Data Store
Amy Tai, **Michael Wei**, Michael J. Freedman, Ittai Abraham, and Dahlia Malkhi
2016 USENIX Annual Technical Conference (USENIX ATC 16), 2016, Denver, CO

HotStorage '16
Silver: A Scalable, Distributed, Multi-versioning, Always Growing (Ag) File System
Michael Wei, Chris Rossbach, Ittai Abraham, Udi Wieder, Steven Swanson, Dahlia Malkhi, and Amy Tai
8th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage 16), 2016, Denver, CO

NVMW '15
AppNVM: A software-defined, application-driven SSD
Matias Bjørling, **Michael Wei**, Jesper Madsen, Javier González, Steven Swanson, and Philippe Bonnet
6th Annual Non-Volatile Memories Workshop 2015 (NVMW 2015), 2015, San Diego, CA

USENIX ATC '14
I/O Speculation for the Microsecond Era
Michael Wei, Matias Bjørling, Philippe Bonnet, and Steven Swanson
2014 USENIX Annual Technical Conference (USENIX ATC 14), 2014, Philadelphia, PA

SOSP '13
Tango: Distributed Data Structures over a Shared Log
Mahesh Balakrishnan, Dahlia Malkhi, Ted Wobber, Ming Wu, Vijayan Prabhakaran, **Michael Wei**, John D. Davis, Sriram Rao, Tao Zou, and Aviad Zuck
Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles, 2013, Farminton, Pennsylvania. DOI: *10.1145/2517349.2522732*

IEEE S&P '13
Welcome to the Entropics: Boot-Time Entropy in Embedded Devices
Keaton Mowery, **Michael Wei**, David Kohlbrenner, Hovav Shacham, and Steven Swanson
2013 IEEE Symposium on Security and Privacy, 2013, Oakland, CA. DOI: *10.1109/SP.2013.46*

SYSTOR '13
Beyond Block I/O: Implementing a Distributed Shared Log in Hardware
Michael Wei, John D. Davis, Ted Wobber, Mahesh Balakrishnan, and Dahlia Malkhi
Proceedings of the 6th International Systems and Storage Conference, 2013, Haifa, Israel. DOI: *10.1145/2485732.2485739*

NSDI '12
CORFU: A Shared Log Design for Flash Clusters
Mahesh Balakrishnan, Dahlia Malkhi, Vijayan Prabhakaran, Ted Wobbler, **Michael Wei**, and John D. Davis
9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12), 2012, San Jose, CA

FAST'11
Reliably Erasing Data from Flash-Based Solid State Drives
Michael Wei, Laura M. Grupp, Frederick E. Spada, and Steven Swanson
Proceedings of the 9th USENIX Conference on File and Stroage Technologies, 2011, San Jose, California

BuildSys'10
Occupancy-Driven Energy Management for Smart Building Automation
Yuvraj Agarwal, Bharathan Balaji, Rajesh Gupta, Jacob Lyles, **Michael Wei**, and Thomas Weng
Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building, 2010, Zurich, Switzerland. DOI: *10.1145/1878431.1878433*

Journal Articles

Leveraging Hyperupcalls To Bridge The Semantic Gap: An Application Perspective
Michael Wei and Nadav Amit
IEEE Data Eng. Bull. *42.1 (2019) pp. 22–35. 2019*

Consistent Clustered Applications with Corfu
Michael Wei, Medhavi Dhawan, Gurprit Johal, Jim Stabile, Vjekoslav Brajkovic, James Chang, Kapil Goyal, Kevin James, Zeeshan Lokhandwala, Anny Martinez Manzanilla, Roger Michoud, Maithem Munshed, Srinivas Neginhal, Konstantin Spirov, Scott Fritchie, Chris Rossbach, Ittai Abraham, and Dahlia Malkhi
SIGOPS Oper. Syst. Rev. *51.1 (Sept. 2017) pp. 78–82. Association for Computing Machinery, 2017,* DOI: *10.1145/3139645.3139658*

CORFU: A Distributed Shared Log
Mahesh Balakrishnan, Dahlia Malkhi, John D. Davis, Vijayan Prabhakaran, **Michael Wei**, and Ted Wobber
ACM Trans. Comput. Syst. *31.4 (Dec. 2013). Association for Computing Machinery, 2013,* DOI: *10.1145/2535930*

Technical Reports

Dynamically scalable, fault-tolerant coordination on a shared logging service
Michael Wei, Mahesh Balakrishnan, John D Davis, Dahlia Malkhi, Vijayan Prabhakaran, and Ted Wobber
Technical Report MSR-TR-2013-40, 2013

SAFE: Fast, Verifiable Sanitization for SSDs
Steven Swanson and **Michael Wei**
Technical Report UCSD TR-cs2011-0963, 2011

Patents

US20230124036A1	In-Place Garbage Collection for State Machine Replication
US20230089659A1	Clock Calibration in a Computing System Using Temperature Sensors
US20220229590A1	Managing Data Lifecycles Through Decay
US11341051B2	Consolidating Shared State for Translation Lookaside Buffer Shootdowns
US20220083476A1	Early acknowledgement of translation lookaside buffer shootdowns
US10545742B2	Annotation-driven framework for generating state machine updates
US20210026785A1	Enforcing code integrity using a trusted computing base
US11055184B2	In-place garbage collection of a sharded, replicated distributed state machine
US11003471B2	Just-in-time hardware for field programmable gate arrays
US11003472B2	Just-in-time hardware for field programmable gate arrays
US10997338B2	Just-in-time hardware for field programmable gate arrays
US10990730B2	Just-in-time hardware for field programmable gate arrays
US10990730B2	Just-in-time hardware for field programmable gate arrays
US11068422B1	Software-controlled interrupts for I/O devices
US10908912B1	Target injection safe method for dynamically inlining branch predictions
US10599835B2	32-bit address space containment to secure processes from speculative rogue cache loads
US10878085B2	Compilation-time checks to secure processes from speculative rogue cache loads
US20190243966A1	Dynamic binary translation to secure processes from speculative rogue cache loads
US10824717B2	Dynamic binary translation to secure processes from speculative rogue cache loads
US10649981B2	Direct access to object state in a shared log
US10642792B2	Distributed transaction conflict resolution
US10819611B2	Dynamic timeout-based fault detection
US10706005B2	File system interface for remote direct memory access
US10635541B2	Fine-grained conflict resolution in a shared log
US10877881B2	In-place garbage collection of a sharded, replicated distributed state machine
US10574571B2	Method for transparently logging and replicating distributed services
US11016767B2	Target injection safe method for inlining instance-dependent calls
US10834255B1	Target injection safe method for inlining large call tables
US10871974B1	Target injection safe method for inlining registration calls
US20190243990A1	Separate cores to secure processes from speculative rogue cache loads

US20190065327A1 Efficient versioned object management

US20190129982A1 Just-in-time multi-indexed tables in a shared log

US20190243776A1 32-bit address space containment to secure processes from speculative rogue cache loads

US20190243965A1 Compilation-time checks to secure processes from speculative rogue cache loads

US20190236229A1 Just-in-time hardware for field programmable gate arrays

US20190235892A1 Just-in-time hardware for field programmable gate arrays

US20190236231A1 Just-in-time hardware for field programmable gate arrays

US20190235893A1 Just-in-time hardware for field programmable gate arrays

US20190236230A1 Just-in-time hardware for field programmable gate arrays

US20180276234A1 Distributed transaction conflict resolution

US20180321963A1 Safe execution of virtual machine callbacks in a hypervisor

service

2024	**PC Member**	ASPLOS
2023	**PC Member**	OSDI
2022–2023	**PC Member**	NSDI
2020	**PC Member**	HotCloud
2018–2019	**PC Member**	USENIX ATC
2018	**PC Member**	ACM SoCC
2017	**PC Member**	ACM Middleware
2017–2024	**PC Member**	HotStorage
2024	**PC Member**	SYSTOR
2011–2015	**Web Chair**	NVMW