Huaicheng LI

Assistant Professor of Computer Science, Virginia Tech | *Systems research at the hardware-software boundary*

Contact

Phone: (540) 231-4482
Gilbert Place 4109, Blacksburg, VA 24060
Google Scholar

Email: huaicheng@vt.edu
Website: huaicheng.github.io
MoatLab: moatlab.github.io

Research Interests

Vision: Hardware innovation outpaces software's ability to harness it. My research unlocks the potential of emerging hardware through systems that bridge innovation and impact. Pond [C8] informs Azure's CXL architecture; FEMU [C13] enables storage research at 100+ institutions.

Areas: Operating Systems, Memory/Storage Systems, Computer Architecture, Datacenter Systems
Current Focus: Building systems for new hardware to improve performance, efficiency, and programmability.

1. *Memory Tiering & Pooling:* Rethinking CXL memory management for datacenter efficiency [C1, C5, C8]
2. *Performance Abstractions:* Designing characterization and prediction models to guide optimizations [C2, C6]
3. *Research Platforms:* Building open-source tools and emulators that enable systems research [C3, C4, C13]
4. *Programmability:* Designing programming models for computational and disaggregated systems [C8, C12]

Research Highlights

Recognition: NSF CAREER Award (2024), Google Research Scholar (2025), 5 Best Paper Awards/Nominations (ASPLOS, FAST, SYSTOR), IEEE Micro Top Picks Honorable Mention (2024)

Publications (since 2022): Pioneering CXL research in characterization, modeling, and tiering/pooling with novel system designs (1× OSDI, 5× ASPLOS, 2× FAST); 7 as corresponding/leading faculty author

Funding: $2.3M total ($1.3M personal share): NSF CAREER ($677K) plus additional NSF and industry grants

Open Source: FEMU (500+ Github stars, widely used at top-venues); Pond (200+ stars, standard CXL simulation approach, foundational for Azure's CXL infrastructure); and research artifacts for reproducible science

Academic Positions

Virginia Tech — Blacksburg, VA
Assistant Professor, Department of Computer Science — 2022–Present

Carnegie Mellon University — Pittsburgh, PA
Postdoctoral Researcher, Parallel Data Lab (PDL) — 2020–2022
Supervisor: Gregory R. Ganger

Education

University of Chicago — Chicago, IL
Ph.D. in Computer Science (*M.S. conferred in 2018*) — 2015–2020
Advisor: Haryadi S. Gunawi
Thesis: Evolving Storage Stack for Predictability and Efficiency [D1]

Wuhan University — Wuhan, China
M.S. in Computer Science (*dropped out to attend the Ph.D. program*) — 2013–2015
B.S. in Computer Science and Technology — 2009–2013

Honors & Awards

ASPLOS '26 Best Paper Award Honorable Mention: PACT [C1]	2026
Google Research Scholar Award	2025
Rising Star Faculty Award, VT CS	2024
NSF CAREER Award	2024
IEEE Micro Top Picks 2024 Honorable Mention: Pond [C8]	2024
ASPLOS '23 Distinguished Paper Award: Pond [C8]	2023
SYSTOR '22 Best Paper Award: Fantastic SSD Internals [C10]	2022
SYSTOR '21 Distinguished Reviewer Award	2021
FAST '18 Best Paper Nominee: Fail-Slow at Scale [C14]	2018
FAST '17 Best Paper Nominee: Tiny-Tail Flash [C16]	2017

Conference Publications ([†] corresponding faculty author; [S] advised student; [M] mentored student)

ASPLOS '26 [C1] Hamid Hadian[S], Jinshu Liu[S⋆], Hanchen Xu[S⋆], Hansen Idden[S], Huaicheng Li[†]. *PACT: A Criticality-First Design for Tiered Memory*. In the 31st ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2026.
⋆ Equal contribution
Best Paper Award Honorable Mention

ASPLOS '26 [C2] Jinshu Liu[S], Hanchen Xu[S], Daniel S. Berger, Marcos K. Aguilera, Huaicheng Li[†]. *Performance Predictability in Heterogeneous Memory*. In the 31st ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2026.

FAST '26 [C3] Dongha Yoon[M⋆], Hansen Idden[S⋆], Jinshu Liu[S], Berkay Inceisci[S], Sam H. Noh, Huaicheng Li[†]. *Cylon: Fast and Accurate Full-System Emulation of CXL-SSDs*. In the 24th USENIX Conference on File and Storage Technologies (FAST), 2026.
⋆ Co-lead authors

FAST '26 [C4] Inho Song[S], Shoaib Qazi[S], Javier González, Matias Bjørling, Sam H. Noh, Huaicheng Li[†]. *Characterizing and Emulating FDP SSDs with WARP*. In the 24th USENIX Conference on File and Storage Technologies (FAST), 2026.

OSDI '25 [C5] Jinshu Liu[S], Hamid Hadian[S], Hanchen Xu[S], Huaicheng Li[†]. *Tiered Memory Management Beyond Hotness*. In the 19th USENIX Symposium on Operating Systems Design and Implementation (OSDI), 2025.

ASPLOS '25 [C6] Jinshu Liu[S], Hamid Hadian[S], Yuyue Wang[S], Daniel S. Berger, Marie Nguyen, Xun Jian, Sam H. Noh, Huaicheng Li[†]. *Systematic CXL Memory Characterization and Performance Analysis at Scale*. In the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2025.

PPoPP '25 [C7] Yongkang Zhang, Haoxuan Yu, Chenxia Han, Cheng Wang, Baotong Lu, Zhifeng Jiang, Yang Li, Xiaowen Chu, Huaicheng Li. *SGDRC: Software-Defined Dynamic Resource Control for Concurrent DNN Inference on NVIDIA GPUs*. In the 30th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), 2025.

ASPLOS '23 [C8] Huaicheng Li[†], Daniel S. Berger, Stanko Novakovic, Lisa Hsu, Daniel Ernst, Pantea Zardoshti, Monish Shah, Ishwar Agarwal, Mark D. Hill, Marcus Fontoura, Ricardo Bianchini. *Pond: CXL-Based Memory Pooling Systems for Cloud Platforms*. In the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2023.
Distinguished Paper Award
IEEE Micro Top Picks 2024 Honorable Mention
Foundational work for Microsoft Azure CXL memory VMs (2025 preview)

ASPLOS '23 [C9] Thomas Kim, Jekyeom Jeon, Nikhil Arora, Huaicheng Li, Michael Kaminsky, David G. Andersen, Gregory R. Ganger, George Amvrosiadis, Matias Bjørling. *RAIZN: Redundant Array of Independent Zoned Namespaces*. In the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2023.

SYSTOR '22 [C10] Nanqinqin Li, Mingzhe Hao, Huaicheng Li, Xing Lin, Tim Emami, Haryadi S. Gunawi. *Fantastic SSD Internals and How to Learn and Use Them.* The 15th ACM International Systems and Storage Conference (SYSTOR), 2022.
Best Paper Award

SOSP '21 [C11] Huaicheng Li, Martin L. Putra, Ronald Shi, Xing Lin, Gregory R. Ganger, Haryadi S. Gunawi. *IODA: A Host/Device Co-Design for Strong Predictability Contract on Modern Flash Storage.* In the 28th Symposium on Operating Systems Principles (SOSP), 2021.

ASPLOS '20 [C12] Huaicheng Li, Mingzhe Hao, Stanko Novakovic, Vaibhav Gogte, Sriram Govindan, Dan R. K. Ports, Irene Zhang, Ricardo Bianchini, Haryadi S. Gunawi, Anirudh Badam. *LeapIO: Efficient and Portable Virtual NVMe Storage on ARM SoCs.* In the 25th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2020.
Deployed in Microsoft Azure production datacenters

FAST '18 [C13] Huaicheng Li, Mingzhe Hao, Michael Hao Tong, Swaminatahan Sundararaman, Matias Bjørling, Haryadi S. Gunawi. *The CASE of FEMU: Cheap, Accurate, Scalable and Extensible Flash Emulator.* In the 16th USENIX Conference on File and Storage Technologies (FAST), 2018.
500+ GitHub stars · Widely used by top-venue papers at OSDI, SOSP, ASPLOS, FAST

FAST '18 [C14] Haryadi S. Gunawi, Riza Suminto, Russell Sears, Casey Golliher, Swaminatahan Sundararaman, Xing Lin, Tim Emami, Weiguang Sheng, Nematollah Bidokhti, Caitie McCaffrey, Gary Grider, Parks M. Fields, Kevin Harms, Robert B. Ross, Andree Jacobson, Robert Riccio, Kirk Webb, Peter Alvaro, H. Birali Runesh, Mingzhe Hao, Huaicheng Li. *Fail-Slow at Scale: Evidence of Hardware Performance Faults in Large Production Systems.* In the 16th USENIX Conference on File and Storage Technologies (FAST), 2018.
Best Paper Nominee
Fast-tracked to ACM TOS; Invited to USENIX ;login:

SOSP '17 [C15] Mingzhe Hao, Huaicheng Li, Michael Hao Tong, Chrisma Pakha, Riza Suminto, Cesar A. Stuardo, Andrew A. Chien, Haryadi S. Gunawi. *MittOS: Supporting Millisecond Tail Tolerance with Fast Rejecting SLO-Aware OS Interface.* In the 26th Symposium on Operating Systems Principles (SOSP), 2017.

FAST '17 [C16] Shiqin Yan, Huaicheng Li, Mingzhe Hao, Michael Hao Tong, Swaminatahan Sundararaman, Andrew A. Chien, Haryadi S. Gunawi. *Tiny-Tail Flash: Near-Perfect Elimination of Garbage Collection Tail Latencies in NAND SSDs.* In the 15th USENIX Conference on File and Storage Technologies (FAST), 2017.
Best Paper Nominee
Fast-tracked to ACM TOS

Journal Publications

TODAES '24 [J1] Ping-Xiang Chen, Dongjoo Seo, Changhoon Sung, Jongheum Park, Minchul Lee, Huaicheng Li, Matias Bjørling, Nikil Dutt. *ZoneTrace: A Zone Monitoring Tool for F2FS on ZNS SSDs*. ACM Transactions on Design Automation of Electronic Systems (TODAES), 2024.

IEEE Micro '23 [J2] Daniel S. Berger, Daniel Ernst, Huaicheng Li, Pantea Zardoshti, Monish Shah, Samir Rajadnya, Scott Lee, Lisa Hsu, Ishwar Agarwal, Mark D. Hill, Ricardo Bianchini. *Design Tradeoffs in CXL-Based Memory Pools for Cloud Platforms.* IEEE Micro Special Issue on Emerging System Interconnects, 2023.

TOS '23 [J3] Huaicheng Li, Martin L. Putra, Ronald Shi, Fadhil I. Kurnia, Xing Lin, Jaeyoung Do, Achmad I. Kistijantoro, Gregory R. Ganger, Haryadi S. Gunawi. *Extending and Programming the NVMe I/O Determinism Interface for Flash Arrays.* ACM Transactions on Storage (TOS), Volume 19, Issue 1, February 2023. [Extended version of C11]

TOS '18 [J4] Haryadi S. Gunawi, Riza Suminto, Russell Sears, Casey Golliher, Swaminatahan Sundararaman, Xing Lin, Tim Emami, Weiguang Sheng, Nematollah Bidokhti, Caitie McCaffrey, Gary Grider, Parks M. Fields, Kevin Harms, Robert B. Ross, Andree Jacobson, Robert Riccio, Kirk Webb, Peter Alvaro, H. Birali Runesh, Mingzhe Hao, Huaicheng Li. *Fail-Slow at Scale: Evidence of Hardware Performance Faults in Large Production Systems.* ACM Transactions on Storage (TOS), Volume 14, Issue 3, November 2018. [Extended version of C14]
Fast-tracked

TOS '17 [J5] Shiqin Yan, Huaicheng Li, Mingzhe Hao, Michael Hao Tong, Swaminatahan Sundararaman, Andrew A. Chien, Haryadi S. Gunawi. *Tiny-Tail Flash: Near-Perfect Elimination of Garbage Collection Tail Latencies in NAND SSDs.* ACM Transactions on Storage (TOS), Volume 13, Issue 3, October 2017. [Extended version of C16]
Fast-tracked

Workshop Publications

HotStorage '23 [W1] Dongjoo Seo, Ping-Xiang Chen, Huaicheng Li, Matias Bjørling, Nikil Dutt. *Is Garbage Collection Overhead Gone? Case study of F2FS on ZNS SSDs.* In the Proceedings of the 15th ACM Workshop on Hot Topics in Storage and File Systems (HotStorage), 2023.

NVMW '23 [W2] Huaicheng Li, Daniel S. Berger, Lisa Hsu, Daniel Ernst, Pantea Zardoshti, Stanko Novakovic, Monish Shah, Samir Rajadnya, Scott Lee, Ishwar Agarwal, Mark D. Hill, Marcus Fontoura, Ricardo Bianchini. *Pond: The Case of CXL Memory Pooling for Cloud Datacenters.* NVMW '23.

Dissertation

PhD '20 [D1] Huaicheng Li, Ph.D. Thesis, *Evolving Storage Stack for Predictability and Efficiency.* University of Chicago. 2020.

MS '18 [D2] Huaicheng Li, M.S. Thesis, *FEMU: Fast, Accurate and Extensible Flash Emulator.* University of Chicago. 2018.

Work Experience

Microsoft Research (Redmond), *Systems Research Group* — Summer 2020
Research Intern working on resource disaggregation for datacenter deployment [ASPLOS '23]

NetApp, *Advanced Technology Group (ATG)* — Spring 2020
Research Intern working on new file system designs for emerging storage hardware

Microsoft Research (Redmond), *Database Group* — Summer 2019
Research Intern working on programmable storage

Microsoft Research (Redmond), *Systems Research Group* — Summer 2018
Research Intern working on offloading cloud storage stack to ARM SoCs [ASPLOS '20]

Service

Leadership Roles

FAST Work-in-Progress **Co-chair**	2027
FAST Artifact Evaluation Committee **Co-chair**	2024, 2025
HotStorage '25 Publicity **Co-chair**	2025
PhD Qualification Exam Committee **Chair**, VT CS	2025

Program Committee – Top Systems Venues

ASPLOS	2023, 2024, 2025, 2026, 2027
FAST	2025, 2026, 2027
NSDI	2024, 2025, 2026
EuroSys	2026, 2027
MICRO	2025
HPCA	2026

External Review Committee (ERC)
ISCA '24, ASPLOS '23 (*Spring/Summer cycles*)

Program Committee – Other Venues
DIMES '26, ICDCS '25, CCGRID '24, SYSTOR '21/'23, APSys '21, NVMW '23, SySDW '23

Proposal Review

NSF Panel	2024, 2025
NSF Ad-hoc, DOE	2024

Journal Reviewer
TOS (2018, 2023, 2024), TOCS (2019, 2023), IEEE Micro (2022), TACO (2024), TC (2019–2023), TPDS (2019), CACM (2018)

Shadow Program Committee

EuroSys '20, EuroSys '18

Artifact Evaluation Committee (AEC)
SOSP '21

External Reviewer
FAST '19, ATC '18

Departmental Service

PhD Qualification Exam Committee, VT CS	2024, 2025 (chair)
Graduate Admission Committee, VT CS	2023–2025
Graduate Student Ministry - Minister for Faculty Hiring, UChicago	2019

Other Service
Session Chair: ASPLOS '26, ASPLOS '23, SYSTOR '21
Mentoring: FAST '24, ASPLOS '23, SOSP '21
BoF Organizer: FAST '24, ATC/OSDI '23 (Reproducibility)

VTURCS Research Symposium Judge	2023
Chameleon Cloud Testbed Student Ambassador	2020

PC Meeting Scribe: HotStorage '20, ATC '18

Talks

Systematic CXL Memory Characterization and Performance Analysis at Scale
Invited Talk, OCP Composable Memory Systems (CMS) [w/ my student, Jinshu Liu (presenter)] 2025

Dissecting Memory Performance: Quantification, Analysis, and Optimization
Invited Talk, Samsung Memory Solutions Lab (MSL) 2024

Pond: CXL-based Memory Pooling for Cloud Platforms

Conference Talk, ASPLOS '23, Vancouver, BC, Canada	2023
Conference Talk, NVMW '23, San Diego, CA, USA	2023
Cornell Networked System Seminar	2023
Guest Lecture at University of British Columbia	2024
Guest Lecture at University of Chicago	2024

Towards Predictable and Efficient Datacenter Storage
Invited Talk, Intel/VMware Crossroads 3D-FPGA Academic Research Center 2022

IODA: Host/Device Co-Design for Strong Predictability Contract on Modern Flash Storage

Conference Talk, SOSP '21, Online	2021
Parallel Data Lab (PDL) Seminar, Carnegie Mellon University, PA, USA	2021

Towards Hardware-based Memory Disaggregation
Invited Talk, Microsoft Research (Redmond), WA, USA 2020

NVMeFS: SmartNIC-centric File System Offloading
Invited Talk, NetApp, CA, USA 2020

Evolving Storage Stack for Predictability and Efficiency

Invited Talk, University of Illinois at Urbana-Champaign, IL, USA	2021
Ph.D. Thesis Defense, University of Chicago, IL, USA	2020
Invited Talk, Carnegie Mellon University, PA, USA	2020
Invited Talk, Microsoft Research - Cambridge, UK	2020
Invited Talk, University of California - Berkeley, CA, USA	2020
Invited Talk, University of Wisconsin - Madison, WI, USA	2020

LeapIO: Efficient and Portable Virtual NVMe Storage on ARM SoCs
Invited Talk, CAS ICT Young Scholar Forum, Beijing, China 2020

Conference Talk, ASPLOS '20, Lausanne, Switzerland 2020
Ph.D. Thesis Proposal, University of Chicago, IL, USA 2019
Invited Talk, Microsoft Research (Redmond), WA, USA 2018

The CASE of FEMU: Cheap, Accurate, Scalable and Extensible Flash Emulator
Master Thesis Defense, University of Chicago, IL, USA 2018
Conference Talk, FAST '18, Oakland, CA, USA 2018

Teaching

Instructor
CS5264/CS4224: (Advanced) Linux Kernel Programming, Virginia Tech 26S, 25S
CS3214: Computer Systems, Virginia Tech 26S, 24F, 22F
CS5204: Operating Systems, Virginia Tech 23F
CS6204: Advanced Operating Systems, Virginia Tech 23S

Co-Instructor
18-746: Storage Systems (Fall 2021), Carnegie Mellon University 2021
Course website: https://course.ece.cmu.edu/ece746/index.html
↪ Co-teaching with Gregory R. Ganger and George Amvrosiadis, ∼100 students (BS/MS/PhD)
↪ Designing and giving lectures on storage management, file systems, etc.
↪ End-to-end class administration: course website, weekly TA meetings, designing quizzes, grading, etc.

Guest Lecturer
18-746: Storage Systems (Fall 2020), Carnegie Mellon University 2020
↪ Topic: "Ins and Outs of Storage Offloading using ARM SoCs"

Students

Current PhD Students
1. Jinshu Liu (**Google PhD Fellow '25**; 5 papers: 1× OSDI, 3× ASPLOS, 1× FAST [C1, C2, C3, C5, C6]) 2022–
2. Hamid Hadian (3 papers: 1× OSDI, 2× ASPLOS [C1, C5, C6]) 2022–
3. Shoaib Asif Qazi (1 paper: 1× FAST [C4]) 2023–
4. Hansen Idden (2 papers: 1× ASPLOS, 1× FAST [C1, C3]) 2024–
5. Hanchen Xu (3 papers: 1× OSDI, 2× ASPLOS [C1, C2, C5]) 2024–
6. Sijia Li 2024–
7. Berkay Inceisci (1 paper: 1× FAST [C3]) 2025–

Co-advised PhD Students (*)
8. Inho Song* (w/ Sam Noh; 1× FAST [C4]) 2023–
9. Yuyue Wang* (at UCLA; 1× ASPLOS [C6]) 2021–

Current Undergraduate Students
10. Zhenyu Zhang (WHU) 2025–
11. Yi Sun (FDU) 2024–

PhD Alumni
12. Sumit Kumar Monga (VT ECE PhD → Postdoc @ INRIA; co-advised w/ Changwoo Min) 2019–2025

Notable Alumni Placements
Zixu Chen, Aditya Shetty → Google · Ronald Shi → Meta · Jiuzhi Yu → AWS · Sumanth Rao → Snowflake
Martin Putra → UChicago PhD · Fadhil Kurnia → UMass PhD · Edward Halim, Yuanzhuo Yang → UW-Madison PhD
Hansen Idden, Hanchen Xu → VT PhD · Subhalakshmi Selvanathan (VT MS) → Microsoft

Student Awards
Jinshu Liu: Google PhD Fellowship '25, Pratt Fellowship
Hamid Hadian (2x), Hanchen Xu, Subhalakshmi Selvanathan: Pratt Fellowship
Jiuzhi Yu, Sumanth Rao: CMU MCDS Best Science Award '22

Thesis Committees
PhD: Sumit Monga (**co-chair**, '25), Ahmad Hossein Yazdani, Ahmad Khan

MS: Subhalakshmi Selvanathan (**co-chair**)

Grants

Total external funding for VT (as PI and Co-PI): $2,338,048.
Personal share: $1,327,300. (External funding for VT, excluding pending grants)

1. "CSR: Performance Criticality as a New Foundation for Memory Systems" 2026–2030
   Huaicheng Li (Sole PI). NSF. $939,328. (*Pending*).
2. "POSE: Phase I: Toward a Community-Driven FEMU Ecosystem for Next-Generation Storage Systems Research and Innovation" 2026–2027
   Huaicheng Li (Sole PI). NSF. $299,946. (*Pending*).
3. "Revisting Operating System Support for Tiered Memory" 2025
   Huaicheng Li (Sole PI). Google. $60,000. (*Unrestricted research gift*).
4. "Rethinking System Stack for the Load-Store I/O Era" 2024–2029
   Huaicheng Li (Sole PI). NSF CNS-2339901. $676,549.
5. "Converged Memory and Storage Systems" 2024–2025
   Huaicheng Li (Lead PI). Samsung. $250,000.
6. "A Cross-stack Approach to Reduce Memory Carbon for Cloud Data Centers" 2023–2026
   Huaicheng Li (Co-PI). NSF CNS-2312785. $1,000,000.
7. "Near-data Processing for Machine Learning Workloads Acceleration" 2023–2024
   Huaicheng Li (Sole PI). 4-VA. $30,000.
8. "CXL for Reduced Memory Management Tax" 2023–2024
   Huaicheng Li (Lead PI). Samsung. $270,151.
9. "Disaggregation and Offloading for Improved System Efficiency" 2023–2024
   Huaicheng Li (Sole PI). The Indonesian Ministry of Education, Culture, Research and Technology. $30,000.
10. "Characterization Driven Data Placement Optimizations for CXL Memory" 2023–2024
    Huaicheng Li (Lead PI). Samsung. $111,348. (*Unrestricted research gift*).
11. "Enhancing Storage Stack Design for the Computational Storage Era" 2023–2024
    Huaicheng Li (Co-PI, **major proposal writer**). Samsung. ∼$120K.
12. New Faculty Mentoring Grant 2023–2024
    Huaicheng Li (PI). Virginia Tech. $1,500.
13. "Fortified Computational Storage Stack for Efficient Application Offloading" 2022–2023
    Huaicheng Li (Co-PI, **major proposal writer**). Samsung. $110K.

Software

1. **SoarAlto:** https://github.com/MoatLab/SoarAlto 2025
   Tiered memory management allocation (Soar) and migration (Alto) policies beyond hotness.
2. **Melody:** https://github.com/MoatLab/Melody 2025
   A systematic CXL performance characterization and analysis framework.
3. **RAZIN:** https://github.com/ZonedStorage/RAIZN-release 2023
   An array of independent zoned namespace SSDs built on top of a virtual zone interface.
4. **Pond:** https://github.com/MoatLab/Pond 2022
   A CXL memory emulator utilizing zero-core NUMA nodes with benchmarking results for more than 100 workloads.
5. **Queenie:** https://github.com/ucare-uchicago/Queenie 2022
   A user-level tool for extracting SSD internal properties.
6. **IODA:** https://github.com/MoatLab/IODA 2021
   A host/device co-designed flash array for strong deterministic performance.
7. **LeapIO:** https://github.com/MoatLab/LeapIO 2020
   A cost-efficient cloud storage stack design that has been deployed in Microsoft datacenters.
8. **FEMU:** https://github.com/MoatLab/FEMU 2018
   A popular storage research platform widely used by top venue papers at ASPLOS, FAST, OSDI, and SOSP, etc.

9. **MittSSD:** https://github.com/ucare-uchicago/mittssd 2018
   An OS design with millisecond service level agreement interface.
10. **ttFlash**: https://github.com/ucare-uchicago/tinyTailFlash 2017
    An SSD architecture design eliminating garbage collection overhead for tiny-tail latencies.
11. Linux Kernel Contributor: Linux Open-Channel SSD Subsystem - pblk (120☆) 2017

Media Coverage

Pond: CXL-Based Memory Pooling Systems for Cloud Platforms [ASPLOS '23]

Software Engineering, https://semiengineering.com 2023
The Next Platform, https://nextplatform.com 2022
Semi Analysis, https://semianalysis.com 2022
Tech Powerup, https://www.techpowerup.com 2022
Screen Hacker, https://www.screenhacker.com 2022

Fail-Slow at Scale [FAST '18]

The Morning Paper, https://blog.acolyer.org, search "fail slow at scale" 2018
ZDNet, https://www.zdnet.com/article/how-clouds-fail-slow 2018
Hacker News, https://news.ycombinator.com/item?id=16463714 2018