

RAMNATHAN ALAGAPPAN

POSTDOCTORAL RESEARCHER
VMWARE RESEARCH GROUP

Curriculum Vitae - October 18, 2021

ADDRESS 4801 Sheboygan Ave, # 715
Madison, WI 53705

WEBSITE <http://pages.cs.wisc.edu/~ra/>
EMAIL ralagappan@vmware.com

RESEARCH INTERESTS

File and Storage Systems, Distributed Systems, and Operating Systems.

Research Summary: *My research improves the reliability and performance of distributed storage systems by co-designing distributed protocols and local-storage stacks.* Distributed storage systems treat local-storage layers as a black box. While this abstraction eases development, it masks vital information about the below layers to distributed protocols, resulting in poor reliability and missed performance opportunities. In my work, I build new distributed systems that use cross-layer information to improve reliability and performance.

My research vision is to make distributed storage systems *future-infrastructure-proof*, i.e., developers can build a storage system for today's target stacks but be confident that the system will function correctly and deliver peak performance on any future infrastructure (e.g., managed environments such as Kubernetes, and rack-scale computers).

EDUCATION

Ph.D. in Computer Sciences

University of Wisconsin – Madison 2019
Advisors: Andrea C. Arpaci-Dusseau and Remzi H. Arpaci-Dusseau
Thesis: Protocol- and Situation-Aware Distributed Storage Systems

M.S. in Computer Sciences

University of Wisconsin – Madison 2018

B.Tech in Information Technology

Coimbatore Institute of Technology, Anna University 2010

HONORS & AWARDS

Distinguished Reviewer	HotStorage 2021
Best Paper Award (CAD)	FAST 2020
MS Azure Credits Research Award for \$50,000	2019
Facebook Distributed Systems Research Award for \$50,000	2019
UW CS Graduate Student Research Award - Best Thesis - Honorable Mention	2019
Best Shadow PC Reviewer	EuroSys 2019
Best Paper Award (Par/Ctrl)	FAST 2018
Best Paper Award (CCFS)	FAST 2017
Best Paper Nominee (Redundancy Does Not Imply Fault Tolerance)	FAST 2017
CS Alumni Scholarship, University of Wisconsin – Madison	2013

PEER-REVIEWED CONFERENCE PUBLICATIONS

- SOSP '21 C12. Aishwarya Ganesan, **Ramnatthan Alagappan**, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Exploiting Nil-Externality for Fast Replicated Storage*. In Proceedings of the 28th ACM Symposium on Operating Systems Principles, 2021.
- FAST '21 C11. Kan Wu, Zhihan Guo, Guanzhou Hu, Kaiwei Tu, **Ramnatthan Alagappan**, Rathijit Sen, Kwanghyun Park, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *The Storage Hierarchy is Not a Hierarchy: Optimizing Caching on Modern Storage Devices with Orthus*. In Proceedings of the 19th USENIX Conference on File and Storage Technologies, 2021.
- OSDI '20 C10. Yifan Dai, Yien Xu, Aishwarya Ganesan, **Ramnatthan Alagappan**, Brian Kroth, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *From Wisckey to Bourbon: A Learned Index for Log-structured Merge Trees*. In Proceedings of the 14th USENIX Conference on Operating Systems Design and Implementation, 2020.
- ATC '20 C09. Anthony Rebello, Yuvraj Patel, **Ramnatthan Alagappan**, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Can Applications Recover from Fsync Failures?* In Proceedings of the 2020 USENIX Annual Technical Conference, 2020.
Fast-tracked to Transactions on Storage
- FAST '20 C08. Aishwarya Ganesan, **Ramnatthan Alagappan**, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Strong and Efficient Consistency with Consistency-aware Durability*. In Proceedings of the 18th USENIX Conference on File and Storage Technologies, 2020.
Best Paper Award
Fast-tracked to Transactions on Storage
- OSDI '18 C07. **Ramnatthan Alagappan**, Aishwarya Ganesan, Jing Liu, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Fault Tolerance, Fast and Slow: Exploiting Failure Asynchrony in Distributed Systems*. In Proceedings of the 13th USENIX Conference on Operating Systems Design and Implementation, 2018.
- FAST '18 C06. **Ramnatthan Alagappan**, Aishwarya Ganesan, Eric Lee, Aws Albarghouthi, Vijay Chidambaram, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Protocol-Aware Recovery for Consensus-Based Storage*. In Proceedings of the 16th USENIX Conference on File and Storage Technologies, 2018.
Best Paper Award
Fast-tracked to Transactions on Storage
Invited to ATC 19 Best of the Rest
- EUROSys'17 C05. Amir Saman Memaripour, Anirudh Badam, Amar Phanishayee, Yanqi Zhou, **Ramnatthan Alagappan**, Karin Strauss, Steven Swanson. *Atomic In-Place Updates for Non-Volatile Main Memories with KaminoTx*. In Proceedings of the European Conference on Computer Systems, 2017.
- FAST '17 C04. Aishwarya Ganesan, **Ramnatthan Alagappan**, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Redundancy Does Not Imply Fault Tolerance: Analysis of Distributed Storage Reactions to Single Errors and Corruptions*. In Proceedings of the 15th USENIX Conference on File and Storage Technologies, 2017.
Best Paper Nominee
Invited to Usenix ;login:
Fast-tracked to Transactions on Storage

- FAST '17 C03. Thanumalayan Sankaranarayana Pillai, **Ramnatthan Alagappan**, Lanyue Lu, Vijay Chidambaram, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Application Crash Consistency and Performance with C2FS*. In Proceedings of the 15th USENIX Conference on File and Storage Technologies, 2017.
Best Paper Award
Fast-tracked to Transactions on Storage
Invited to ATC 18 Best of the Rest
- OSDI '16 C02. **Ramnatthan Alagappan**, Aishwarya Ganesan, Yuvraj Patel, Thanumalayan Sankaranarayana Pillai, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Correlated Crash Vulnerabilities*. In Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation, 2016.
- OSDI '14 C01. Thanumalayan Sankaranarayana Pillai, Vijay Chidambaram, **Ramnatthan Alagappan**, Samer Al Kiswany, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *All File Systems Are Not Created Equal: On the Complexity of Crafting Crash-Consistent Applications*. In Proceedings of the 11th USENIX Conference on Operating Systems Design and Implementation, 2014.
Invited to Communications of the ACM
Invited to ACM Queue

PEER-REVIEWED WORKSHOP PUBLICATIONS

- HotOS '21 W04. Xudong Sun, Lalith Suresh, Aishwarya Ganesan, **Ramnatthan Alagappan**, Michael Gasch, Lilia Tang, Tianyin Xu. *Reasoning about Modern Datacenter Infrastructures using Partial Histories* 18h Workshop on Hot Topics in Operating Systems, 2021.
- NVMW '21 W03. Kan Wu, Zhihan Guo, Guanzhou Hu, Kaiwei Tu, **Ramnatthan Alagappan**, Rathijit Sen, Kwanghyun Park, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *The Storage Hierarchy is Not a Hierarchy: Optimizing Caching on Modern Storage Devices with Orthus* Non-volatile Memory Workshop, 2021.
- HotStorage '20 W02. Konstantinos Kanellis, **Ramnatthan Alagappan**, Shivaram Venkataraman. *Too Many Knobs to Tune? Towards Faster Database Tuning by Pre-selecting Important Knobs*. 12th Workshop on Hot Topics in Storage and File Systems, 2020.
- HotOs '15 W01. **Ramnatthan Alagappan**, Vijay Chidambaram, Thanumalayan Sankaranarayana Pillai, Aws Albarghouthi, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Beyond Storage APIs: Provable Semantics for Storage Stacks*. 15th Workshop on Hot Topics in Operating Systems, 2015.

PEER-REVIEWED JOURNAL PUBLICATIONS

- TOS '21 J05. Anthony Rebello, Yuvraj Patel, **Ramnatthan Alagappan**, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Can Applications Recover from fsync Failures?* ACM Transactions on Storage (TOS), June 2021. **Fast-tracked**
- TOS '21 J04. Aishwarya Ganesan, **Ramnatthan Alagappan**, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Strong and Efficient Consistency with Consistency-aware Durability*. ACM Transactions on Storage (TOS), January 2021. **Fast-tracked**
- TOS '18 J03. **Ramnatthan Alagappan**, Aishwarya Ganesan, Eric Lee, Aws Albarghouthi, Vijay Chidambaram, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Protocol-Aware Recovery for Consensus-Based Distributed Storage*. ACM Transactions on Storage (TOS), October 2018.
Fast-tracked

- TOS '17 J02. Aishwarya Ganesan, **Ramnatthan Alagappan**, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Redundancy Does Not Imply Fault Tolerance: Analysis of Distributed Storage Reactions to File-System Faults*. ACM Transactions on Storage (TOS), September 2017. **Fast-tracked**
- TOS '17 J01. Thanumalayan Sankaranarayana Pillai, **Ramnatthan Alagappan**, Lanyue Lu, Vijay Chidambaram, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Application Crash Consistency and Performance with C2FS*. ACM Transactions on Storage (TOS), September 2017. **Fast-tracked**

OTHER PUBLICATIONS

- ;login: P04. Aishwarya Ganesan, **Ramnatthan Alagappan**, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Redundancy Does Not Imply Fault Tolerance: Analysis of Distributed Storage Reactions to Single Errors and Corruptions*. ;login: The USENIX Magazine, Summer 2017. **Invited**
- MSR TR P03. Yanqi Zhou, **Ramnatthan Alagappan**, Amir Samam Memaripour, Anirudh Badam, David Wentzlauff. *Hybrid NVM Enabled Datacenter Design and Optimization*. MSR-TR-2017-8, February 2017.
- ACMQueue P02. Thanumalayan Sankaranarayana Pillai, Vijay Chidambaram, **Ramnatthan Alagappan**, Samer Al Kiswany, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Crash Consistency: Rethinking the Fundamental Abstractions of the File System*. ACM Queue, July 2015. **Invited**
- CACM P01. Thanumalayan Sankaranarayana Pillai, Vijay Chidambaram, **Ramnatthan Alagappan**, Samer Al Kiswany, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Crash Consistency*. Communications of the ACM - Vol. 58, No. 10, October 2015. **Invited**

WIP POSTERS

- NVMW '18: Amir Saman Memaripour, Anirudh Badam, Amar Phanishayee, Yanqi Zhou, **Ramnatthan Alagappan**, Karin Strauss, Steven Swanson. *Atomic In-Place Updates for Non-Volatile Main Memories with KaminoTx*.
- FAST '16: Thanumalayan Pillai, **Ramnatthan Alagappan**, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. *Simple Crash Consistency With Streams*.

RESEARCH IMPACT

Corruption-tolerant Replication. The CTRL protocol from my FAST '18 paper has been adopted and implemented in TigerBeetle ([Link1](#), [Link2](#)), a financial database, making it resilient to storage corruptions and errors.

ErrFS and ErrBench. ErrFS is a user-level FUSE file system that systematically injects file-system faults. Ideas from ErrFS have been adopted by other popular testing tools. ErrBench is a suite of distributed-storage-system workloads which drives systems to interact with their local storage. Through ErrFS and ErrBench, we have exposed many serious bugs in popular distributed systems such as ZooKeeper, Cassandra, and Kafka. [Link to Artifacts](#)

PACE. PACE is a framework to systematically generate and explore persistent states that can occur in a distributed execution, exposing crash vulnerabilities in distributed storage systems. PACE found 26 serious, real-world bugs in popular systems including ZooKeeper, Redis, etcd, and Kafka. Many bugs found by PACE have been fixed by developers. [Link to Artifacts](#)

ALICE. ALICE is a crash-consistency testing framework that I helped build. ALICE has been adopted by others (including an open-source version). ALICE found several real-world bugs in 12 widely used commercial storage software products, including Google's LevelDB, Git, and SQLite.

[Link](#)

PRESS ARTICLES ON RESEARCH

THE MORNING PAPER. Protocol-Aware Recovery for Consensus-Based Storage Link to Article	Feb 2018
ZDNET. Eliminating Storage Failures in the Cloud Link to Article	Feb 2018
THE MORNING PAPER. Crash Consistency and Performance with CCFS Link to Article	Mar 2017
THE MORNING PAPER. Redundancy Does Not Imply Fault Tolerance Link to Article	Mar 2017
DHSR's BLOG. Redundancy Does Not Imply Fault Tolerance Link to Article	Mar 2017
STORAGEMOJO. Redundancy Does Not Imply Fault Tolerance Link to Article	Mar 2017
THE MORNING PAPER. All File Systems are Not Created Equal Link to Article	Feb 2016

TEACHING

Instructor , <i>University of Wisconsin – Madison</i> CS 739 - Distributed Systems (graduate-level) Link to Course Webpage Course evaluation score: 6.42/7.00 (ranked 2nd among Spring 2020 graduate courses) Nominated for SACM CoW Award (yearly teaching award given to UW CS professors) Description: I designed a graduate-level distributed systems course. This course was heavily research oriented: every class, students read 1-2 foundational papers in distributed systems on topics including fault tolerance, consensus (e.g., Paxos), distributed transactions, BFT, and distributed storage. Students also did a considerably large research project. Lecture notes, assignments, and exams are available on the course web page.	SPRING '20
Teaching Assistant , <i>University of Wisconsin – Madison</i> CS 537 - Intro to Operating Systems Responsibilities: Guest lectures, exam review lectures, designing and evaluating assignments on the xv6 research operating system.	SPRING '19
Guest Lectures , <i>University of Wisconsin – Madison</i> Shivaram Venkatramans's CS 537 (concurrency, RAID) Mike Swift's CS 736 (AFS) Remzi Arpaci-Dusseau's CS 739 - (Paxos, storage faults)	SPRING '20, SPRING '19 FALL '18 FALL '18, FALL '17
Teaching Assistant , <i>University of Wisconsin – Madison</i> CS 302 - Introduction To Programming	FALL '13

STUDENT MENTORING

VMware Research

Yi Xu, graduate student at UC San Diego
Research Internship Mentor
Exploiting Persistent Memory in SplinterDB

Undergrads at UW Madison

Neil Perry, now a graduate student at Stanford University
Corruption Analysis of Ethereum Blockchain

Graduate Students at UW Madison

Yifan Dai, Yien Xu
Learned Indexes for Log-Structured Merge Trees (CS 739 final project, OSDI 2020)

Dax Chen, Yi-Shiun Chang, Chia-Wei Chen, Pei-Hsuan Wu
Performance and Reliability Isolation in ZooKeeper

Sreya Dutta Roy, Nikita Kad, Venkat Allam, Shreeshrita Patnaik
Predicted Ordering in Geo-replicated Logs

Akshat Jain, Grishma Gupta, Venkata Malireddy
Learning Based Ordering for Replicated State Machines

Ruohui Wang, Kaiwei Tu, Max (Mengxiao) Zhang, Emma (Yi) He
Read-triggered Durability for HDFS

Aashish Richhariya, Akanksha, Sanchit Jain
Consistency at the Edge

Muthunagappan Muthuraman, Srivatsan Ramesh, Suryadev Sahadevan Rajesh, Vinith Venkatesan
Consistency-Aware Durability for Highly Available Systems

Deepak Srinath, Lokit Kumar Paras, Nithin Venkatesh, Phanindra Moganti
Speculative Geo-Replicated Message Ordering

Kumar Biplav, Aditya Rungta, Nisarg Shah, Shaurya Shekhar
Fast Consensus for Fast Storage

PROFESSIONAL EXPERIENCE

VMware Research Group Palo Alto, CA
Postdoctoral Researcher, VMware Research Group FALL '20 –

Microsoft Research Redmond, WA
Research Intern, Systems Research Group SUMMER '15
Mentor: Anirudh Badam

Microsoft Research Bangalore, India
Research Intern, Mobility, Networks, and Systems Group SUMMER '14
Mentor: Ramachandran Ramjee

Microsoft Hyderabad, India
Software Development Engineer JUL '10 – JUN '13

SERVICE

HotStorage '22 Program Committee	2021
SOSP '21 Ask-Me-Anything Co-chair	2021
OSDI '21 Mentoring	2021
EuroDW '21 Mentoring	2021
Journal of Systems SEB Co-chair	2021
EuroDW '21 Program Committee	2021
HotStorage '21 Program Committee (Distinguished Reviewer)	2021
Systor '21 Program Committee	2021
ACM Transactions on Computer Systems, Reviewer	2020
HotStorage '20 Program Committee	2020
SOSP '19 Artifact Evaluation Committee	2019
Eurosys '19 Shadow PC (Best Reviewer)	2019
ACM Transactions on Storage, Reviewer	2018
FAST '18, External Reviewer	2018
EuroSys '17, Contributor to PC Reviews	2017
OSDI '16, External Reviewer	2016
FAST '16, External Reviewer	2016

PRESENTATIONS & INVITED TALKS

Co-designing Distributed Systems and Storage Stacks University of Waterloo (invited)	OCT '21
Reliable Distributed Storage: A Local-storage Perspective Rutgers University (invited)	AUG '20
Reliable Distributed Storage: A Local-storage Perspective VMware Research Group (postdoc interview talk)	JUN '20
Protocol-Aware Recovery for Consensus-Based Storage Usenix ATC (invited conference talk)	JUL '19
Storage Systems at the Edge NSF-VMWare ECDI Summit (invited)	Nov '18
Fault-Tolerance, Fast and Slow Usenix OSDI (conference talk)	OCT '18
Protocol-Aware Recovery for Consensus-Based Storage SNIA Storage Developer Conference (invited)	SEP '18
Resiliency to Storage Faults in Distributed Systems Google Madison (invited)	MAY '18
Protocol-Aware Recovery for Consensus-Based Storage Usenix FAST (conference talk)	FEB '18
Rethinking Consensus with Local Storage in Mind SCI Labs Kickoff Meeting	MAY '17

Correlated Crash Vulnerabilities

Usenix OSDI (conference talk)

OCT '16

Correlated Crash Vulnerabilities

Microsoft Gray Systems Lab (invited)

JUN '16

GRANTS

Facebook Distributed Systems Research Award for \$50,000 - with Aishwarya Ganesan, Andrea Arpaci-Dusseau, and Remzi Arpaci-Dusseau to work on distributed storage reliability, especially blockchains.

MS Azure Credits Research Award for \$50,000 - with Andrea Arpaci-Dusseau and Remzi Arpaci-Dusseau to work on distributed storage reliability, building new testing and analysis frameworks

Travel grants for OSDI '14, FAST '17, FAST '18

REFERENCES

Available upon request