YOUMIN CHEN

http://storage.cs.tsinghua.edu.cn/~cym
Room 8-201, East Main Building, Tsinghua University, Beijing, 100084, China
chenym16@mails.tsinghua.edu.cn

ABOUT ME

My research area mainly covers distributed systems and storage systems. I'm particularly interested in building (distributed) file systems, key-value stores and transactional systems leveraging emerging devices (e.g., RDMA, programmable hardware). My work has been published on FAST, ASPLOS, Eurosys, USENIX ATC, DAC, VLDB, ICDE, etc.

EDUCATION

- **Ph.D. Tsinghua University** Aug. 2016 - Present
  Computer Science, Storage Research Group.
  Advisor: Prof. Jiwu Shu and Prof. Youyou Lu.
- **Visiting Scholar. University of Wisconsin - Madison** Nov. 2019 - Nov. 2020
  Computer Sciences, The ADvanced Systems Laboratory (ADSL)
  Advisor: Prof. Remzi H. Arpaci-Dusseau, Andrea C. Arpaci-Dusseau and Xiangyao Yu
- **B.Sc. Beihang University** Sept. 2012 - June 2016
  Computer Science.
  ShenYuan Honors College (f.k.a., School of Advanced Engineering)

RESEARCH PUBLICATION

Pre-print

- **Youmin Chen**, Xiangyao Yu, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, and Jiwu Shu. "*High Throughput and Low Tail Latency Transactions with Plor*" (In submission to OSDI'21, CCF-A)

Conference Papers

- Fan Yang **Youmin Chen**, Qing Wang, Youyou Lu, and Jiwu Shu. "*Tolerating Skewed Workloads in Secure In-memory Key-Value Store*" (**ICDE'21, CCF-A**, accepted), 2021
- **Youmin Chen**, Youyou Lu, Bohong Zhu, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, and Jiwu Shu. "*KucoFS: a Kernel/User-level Collaborative File System for Non-Volatile Memories*" (**FAST'21, CCF-A**), 2021
- Qing Wang, Youyou Lu, Junru Li, **Youmin Chen**, Erci Xu, and Jiwu Shu. "*Concordia: Distributed Shared Memory with In-Network Cache Coherence*" (**FAST'21, CCF-A**), 2021.
- **Youmin Chen**, Youyou Lu, Kedong Fang, Qing Wang, and Jiwu Shu. "*µTree: a Persistent B+-Tree with Low Tail Latency*", 46th International Conference on Very Large Data Bases (**VLDB'20**, ***CCF-A***), 2020.
- **Youmin Chen**, Youyou Lu, Fan Yang, Qing Wang, Yang Wang, and Jiwu Shu. "*FlatStore: an Efficient Log-Structured Storage Engine for Persistent Memory Key-value Stores*", Proceedings of the 25th International Conference on Architectural Support for Programming Languages and Operating Systems (**ASPLOS'20**, ***CCF-A***), 2020.
- Fan Yang, Youyou Lu, **Youmin Chen**, and Jiwu Shu. "*No Compromises: Secure NVM with Crash Consistency, Write-Efficiency and High-Performance*", Design Automation Conference (**DAC'19**, ***CCF-A***), 2019.
- **Youmin Chen**, Youyou Lu, and Jiwu Shu. "*Scalable RDMA RPC on Reliable Connection with Efficient Resource Sharing*", European Conference on Computer Systems (**Eurosys'19**, ***TH-CPL-A***), 2019.
- Youyou Lu, Jiwu Shu, **Youmin Chen**, and Tao Li. "*Octopus: an RDMA-enabled Distributed Persistent Memory File System*", USENIX Annual Technical Conference (**USENIX ATC'17**, ***CCF-A***), 2017.

Journal Papers

- Bohong Zhu, **Youmin Chen**, Qing Wang, Youyou Lu, and Jiwu Shu. "*Octopus$^+$: an RDMA-enabled Distributed Persistent Memory FileSystem*", ACM Transactions on Storage (**TOS**, ***CCF-A***, accepted, 2021).

- Jiwu Shu, **Youmin Chen**, Qing Wang, Bohong Zhu, Junru Li, and Youyou Lu. "*TH-DPMS: Design and Implementation of an RDMA-enabled Distributed Persistent Memory Storage System*", ACM Transactions on Storage (**TOS**, ***CCF-A***, accepted, 2020).
- Jiwu Shu, **Youmin Chen**, Qingda Hu, and Youyou Lu. "*The Development of System Software on Non-volatile Main Memory*". SCIENTIA SINICA Informationis, 2020. (***TH-CPL-A***)
- **Youmin Chen**, Youyou Lu, Pei Chen, and Jiwu Shu. "*Efficient and Consistent NVMM Cache for SSD-based File System*", IEEE Transactions on Computers (**TC**, ***CCF-A***), (2018).
- **Youmin Chen**, Jiwu Shu, Jiaxin Ou, and Youyou Lu. "*HiNFS: A persistent memory file system with both buffering and direct-access*", ACM Transactions on Storage (**TOS**, ***CCF-A***) 14.1 (2018): 4.
- **Youmin Chen**, Youyou Lu, Shengmei Luo, and Jiwu Shu. "*Survey on RDMA-based Distributed Storage Systems*", Journal of Computer Research and Development, 2019, 56(2): 227-239. (***TH-CPL-B***, in Chinese).
- **Youmin Chen**, Bohong Zhu, Yinjun Han, Yaofeng Tu, and Jiwu Shu. "*A Hybrid Approach for Managing Data Pages in Persistent Memory File Systems*". Journal of Computer Research and Delevopment (***TH-CPL-B***, in Chinese), 2020.
- Fan Yang, **Youmin Chen**, Haiyu Mao, Youyou Lu, and Jiwu Shu. ShieldNVM: An Efficient and Fast Recoverable System for Secure Non-Volatile Memory. ACM Transactions on Storage (**TOS**, ***CCF-A***), 2020.

Workshop and Posters

- Youyou Lu, Jiwu Shu, **Youmin Chen**, Tao Li. "*Octopus: an RDMA-enabled Distributed Persistent Memory File System*", Non-Volatile Memory Workshop (**NVMW'19**). 2019.
- **Youmin Chen**, Jiwu Shu, Youyou Lu "*Scalable and Reliable RDMA*", Proceedings of the ACM Symposium on Cloud Computing (**SoCC'18**). ACM, 2018.

PROJECTS

Octopus - an RDMA-enabled distributed persistent file system 2015 - 2016
I designed the system with Youyou. I implemented most code.

· Non-volatile memory (NVM) and remote direct memory access (RDMA) provide extremely high performance, but existing distributed file systems leave such high-speed hardware under-utilized.
· Octopus abstracts a shared persistent memory pool and redesigns the file system internal mechanisms. As far as we know, this is the first RDMA-enabled distributed persistent memory file system in the world.
· Published the paper on the top conference USENIX ATC as the first student author and now it has more than 60 citations.

ScaleRPC - improving the scalability of RC-based RDMA RPC 2016 - 2017
I created the idea and designed the system.

· RDMA suffers performance collapse when transferring data to an increasing number of clients on reliable connection because of resource contention in the NIC cache and CPU cache.
· ScaleRPC is an efficient RPC primitive based on RC RDMA to deliver scalable performance. It achieves so by grouping and rating the network connections, so as to balance the saturation and thrashing of the cache space.

FlatStore: Reducing the Flushing Overhead in Persistent Memory Key-Value Stores 2017 - 2018
I created the idea and designed the system, Fan and Qing helped with the code.

· As the emerging storage and network devices become extremely fast, CPU becomes the new bottleneck, especially when processing high concurrent and small- sized requests as in existing key value stores.
· FlatStore organize the small-sized key- value indexes with a log-structured OpLog, and fully exploits the benefits of batching, so as to better utilize the CPU resources.

KucoFS - a persistent memory file system with kernel/user-level cooperation 2018-2019
I created the idea and designed the system. Bohong helped with the code.

· Recent works propose to design persistent memory file system in user-space to avoid the high VFS and syscall overhead in the kernel, but unfortunately leave the file system vulnerable to malicious writes and hard to scale.

· KucoFS is a kernel and user-level collaborative file system to achieve both efficiency of user level designs and data protection of kernel-level ones.

CHRONUS - an in-memory transactional system with high throughput and low tail latency 2019-2020
I created the idea and designed the system.

· Despite many years of research on tail latency in different layers in the operating system, we find that scant attention has been paid to another aspect: *how do request conflicts impact tail latency of these data services?*

· To answer this question, we introduce *pessimistic locking and optimistic reading*, or PLOR, a hybrid concurrency control scheme that enables extremely low tail latency while delivering excellent throughput. The central idea behind PLOR is that it adopts a pessimistic approach to acquire locks before actually accessing a record from the database, but allows transactions to read records without being blocked when lock conflicts occur.

SKILLS

Languages	Mandarin Chinese (Naive), English (Fluent).
Coding	C, C++, LaTeX, JAVA, Go.

MISCELLANEOUS EXPERIENCE

Awards

- **First-class Scholarship, Tsinghua University** 2018/2019
- **Ruian Scholarship, Beihang University** 2015
- **Second Prize Scholarship for Freshman, Beihang University** 2012

Invited Talks

- **Octopus: an RDMA-enabled Distributed Persistent Memory File System**
  Santa Clara, CA, USA 2017
  UC San Diego, CA, USA 2018
- **FlatStore: an Efficient Log-Structured Storage Engine for Persistent Memory Key-value Stores**
  Lausanne, Switzerland 2020