

	<b>姓名:</b>	舒继武
	<b>职称:</b>	IEEE Fellow 国家杰出青年科学基金获得者、长江学者 信息学院院长 教授、博士生导师
	<b>E-mail:</b>	jwshu@xmu.edu.cn
	<b>个人主页:</b>	<a href="http://storage.cs.tsinghua.edu.cn/~jiwu-shu">http://storage.cs.tsinghua.edu.cn/~jiwu-shu</a>
	<b>研究方向:</b>	非易失性内存存储系统、闪存存储系统、分布式存储系统、数据存储可靠性与安全、近数据智能存算融合系统

## 【社会兼职】

- 中国计算机学会：第十二届理事会理事 (2020-);
- 中国计算机学会：信息存储技术专业委员会主任 (2020-);
- 中国大数据智能计算专家委员会：副主任 (2019-2021);
- 中国通信学会：云计算和大数据应用委员会副主任委员 (2018-2021);
- 灾备技术国家工程实验室：副主任、技术委员会委员 (2008-);
- 信息存储系统教育部重点实验室学术委员会委员 (2017-);
- 清华大学信息科学技术学院学术委员会委员 (2011-);
- 863 计划“海量存储系统关键技术”重大项目：总体专家组副组长 (2008-2010);
- 《ACM Transaction on Storage》：Associate Editor (2016-);
- 《IEEE Transactions on Computers》：Associate Editor (2019-);
- 《Journal of Computer Science and Technology》：Member of the Editorial Board (the area of Computer Architecture and Systems) (2019-);
- 《计算机学报》：编委 (2010-);
- 《软件学报》：责任编委 (2012-);
- 《计算机研究与发展》：编委 (2010-);
- 《大数据》：编委 (2015-).

## 【研究领域】

- ① 基于非易失存储器件的新型存储系统与技术
- ② 闪存存储系统与技术
- ③ 网络（/云/大数据）存储系统与关键技术
- ④ 数据存储可靠性与安全
- ⑤ 近数据智能存算融合系统

## 【研究概况】

- ① 在基于非易失存储器件的新型存储系统与技术方面，基于“计算与内存分离”和“功能抽象”的设计思想，提出了分布式持久性共享内存框架，并攻克了其数据结构、内存管理、一致性与安全等方面系列难题，形成了内存、对象、事务、复制等灵活高效的统一接口，研制出了持久性内存文件系统与键值存储系统，构建出了新型分布式持久性内存存储系统 TH-DPMS，打破计算存储紧耦合设计，充分发挥了新型硬件优势；相关主要学术论文发表在 SOSP'21、OSDI'2021、FAST'2021(2)、ICDE'2021、ASPLOS'2020、MICRO'2018/2013、USENIX ATC'2017(2)、EuroSys'2019/2016、VLDB'2020、DAC'2020/2019/2016/2015、DSN'2014、MSST'2020/2016/2015、ICCD'2014 等重要国际会议和 ACM TOS 2020/2018/2016、IEEE TC 2021、IEEE TPDS 2019 等权威国际期刊上，并获得 MSST'2015 国际会议的 Best Papers Nominated。相关成果获得 2020 年华为首届“奥林帕斯奖”和百万悬红，并获得 2020 年“CCF 科学技术奖”技术发明一等奖；该方面研究已获得国家 863 项目、国家重点研发项目等资助外，还获得华为、阿里、快手、中兴等公司资助，部分成果已转化并获得应用。
- ② 在闪存存储系统与技术方面，针对 Flash 的特有特性，提出了软件直管的开放通道（Open-Channel SSD）闪存存储架构与相关技术，获 Linux4.4 内核及后续支持；研制出了闪存文件系统 ParaFS、键值闪存存储系统 FlashKV、分布式闪存存储系统 OCStore 等，构建出了软硬件协同的闪存存储系统 TH-SSS，在性能、可靠性等方面有显著优势；相关主要学术论文发表在 OSDI'2020、FAST'2014/2013、MICRO-54、USENIX ATC'2021/2016、SC'2017、DAC'2020/2019、ICDCS'2020/2019、MSST'2020、IPDPS'2014、ICCD'2013 等重要国际会议和 IEEE TC 2019/2016/2015、IEEE TPDS 2018、ACM TOS 2019 等权威国际期刊上，并获得 NVMSA'2015 国际会议的 Best Papers。相关成果获 2019 年教育部科学技术发明奖一等奖；该方面研究已获得国家 863 项目和三星、华为、中兴等公司资助，部分成果已转化并获得应用，该成果的发明专利获得转让和许可费达 1100 万。
- ③ 在网络（/云/大数据）存储系统与关键技术方面。主要针对网络（/云/大数据）存储系统中面临的性能、扩展性、兼容性、易用性、安全性、易管理等问题进行研究，在存储虚拟化、存储系统在线扩展方法、元数据管理方法、存储 Cache 策略、分级存储、面向 SAN 的并行文件系统、云存储安全等方面的研究取得了一些进展。研制出了国内最早的具有自主知识产权的存储区域网络型存储系统——“清华海量存储网络系统 TH-MSNS”，2006 年前，TH-MSNS 系统已有近 100 套应用在审计署、公安局、电信、油田、高校等行业或部门；同时，也研制了相关的存储子系统，如：存储管理系统 TH-MSIS、分布式虚拟化存储系统 TH-DVSM、带外虚拟化存储系统 TH-VSM、备份系统 TH-Backup、快照系统 TH-Snapshot、远程镜像系统 TH-Mirror、面向 SAN 的并行文件系统 Redbud、分级存储系统 AIP、云存储安全系统 Corslet，一些系统已分别在电信、教育等部门获得实际应用；相关主要学术论文发表在 IEEE TC 2016/2014/2010/2007/2005、ACM TOS 2010/2007、IEEE TPDS 2016/2015/2014、JPDC 2014/2018 等权威国际期刊和 IPDPS'2016、CODES/ISSS'2016、IWQos'2014/2013、ICPP'2012、CLUSTER'2012、ASP-DAC'2013 等重要国际会议上。这些系统与应用成果分别获得 2007 年获国家科技进步二等奖、2015 年国家科学技术发明二等奖等、2014 年中国电子学会科学技术发明一等奖等。
- ④ 在数据存储可靠性方面。针对存储纠删码容错能力（容 1 个、2 个和多个错误）、性能（编码效率）和存储利用率等方面存在的问题进行研究，提出并实现了一些相关的高效实用的存储纠删码技术，提高了存储系统的可靠性，相关主要学术论文发表在 USENIX ATC'2021、DSN'2016/2014、SRDS'2020/2017/2016/2015/2014、IPDPS'2020/2015、

ICPP'2015/2013/2012 等重要国际会议和 IEEE TC 2017/2016/2010、IEEE TPDS 2019/2018/2017/2016、IEEE TDSC 2020、ACM TOS 2012/2009 等权威国际期刊上，并获得 2010 年 IEEE TC 期刊的 Featured Article、SRDS'2015 国际会议的 Nominated Best Papers。

- ⑤ 研制出近数据智能存算融合系统 TH-iSSD，融合了存储与计算的分布式平台，在近数据端通过存算融合的方法以极高的效率完成数据的感知、存储与计算的融合，具有性能高、功耗低、集成度高、可靠性好等特点；已分别基于航天、航空、气象、医疗、工业机器人等多个场景，研制出存算一体化平台 TH-AeroStor、智能边缘存算融合系统 TH-SmartEdge 等。相关发明专利转让和许可费达 220 万。

## 【主要研究项目和课题】

- 国家重点研发计划重点专项项目: TB 级持久性内存存储技术与系统 (2018-2021);
- 国家重点研发计划重点专项课题: 分布式持久性共享内存框架研制及存储系统应用示范 (2018-2021);
- 北京市科委重点项目之课题: 异构混合内存系统的大数据处理软件平台架构与优化研究 (2015-2016);
- 863 项目: 面向大数据的先进存储结构及关键技术(2013-2015);
- 863 课题: 基于新型存储器件的存储体系结构与关键技术研究(2013-2015);
- 国家自然科学基金重点基金: 新型分布式存储系统的高可靠性关键技术研究 (2019-2023);
- 国家自然科学基金重点基金:网络公用数据存储系统可靠性及其灾备研究(2013-2017);
- 国家工信部子课题: 云计算平台上面向移动通信业务云服务关键技术与系统(2013-2015);
- 国家科技支撑计划子课题: 存储虚拟化软件开发及产业化(2010-2013);
- 国家杰出青年科学基金: 大规模网络存储系统的可靠性关键技术研究 (2010-2013);
- 国家自然科学基金: 大规模网络存储系统的容错技术研究 (2009-2011);
- 973 课题: 下一代互联网存储按需部署模型与服务质量研究 (2004-2009);
- 863 重大课题子课题: 海量存储系统软件和关键技术 (2009-2010).

## 【奖励与荣誉】

- 北京市科学技术进步二等奖——高可扩展海量网络存储系统 (2004);
- 高等学校科学技术进步二等奖——深腾 2600 网络服务器系统 (2007);
- 国家科技进步二等奖——高性能集群计算机与海量存储系统 (2007);
- 中国电子学会技术发明一等奖——面向社区共享的高可用云存储系统 (2014);
- 国家科学技术发明二等奖——面向社区共享的高可用云存储系统 (2015);
- 江苏省科学技术奖一等奖——大规模数据服务系统与平台的关键技术及产业应用 (2018);
- 中国产学研合作促进会创新成果奖一等奖——基于 TB 级数据架构的需求设计开发一体化软件工程智能技术研发及行业应用 (2018);
- 教育部科学技术技术发明一等奖——软硬件协同的闪存存储系统关键技术及应用 (2019);
- 中国计算机学会科学技术奖技术发明一等奖——持久性内存存储系统关键技术与应用 (2020).
- 清华大学学术新人奖 (2005);
- 教育部: 新世纪优秀人才 (2005);

- 中创软件人才奖(2007);
- 国家自然科学基金委: 国家杰出青年科学基金获得者 (2009);
- 教育部长江学者特聘教授 (2014);
- IEEE Fellow (2018);
- 中国计算机学会会士 (2018)
- 华为“奥林帕斯奖”(2020).

## 【学术成果】

### ● 代表性会议论文

- [1] Xiaojian Liao, Youyou Lu, Zhe Yang, Jiwu Shu, Crash Consistent Non-Volatile Memory Express, The 28th *ACM Symposium on Operating Systems Principles (SOSP)*, Koblenz, Germany, October 25-28, 2021 (Acceptance rate: 54/348=15.52%)
- [2] Congming Gao, Xin Xin, Youyou Lu, Youtao Zhang, Jun Yang, Jiwu Shu, ParaBit: Processing Parallel Bitwise Operations in NAND Flash Memory based SSDs, the 54th *IEEE/ACM International Symposium on Microarchitecture (MICRO-54)*, October 18-22, 2021, Athens, Greece. (Acceptance rate: 94/430=21.86%)
- [3] Qing Wang, Youyou Lu, Junru Li, Jiwu Shu, Nap: A Black-Box Approach to NUMA-Aware Persistent Memory, the 15th *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 14-16 July 2021 (Acceptance rate: 31/165=18.8%)
- [4] Xiaojian Liao, Youyou Lu, Erci Xu, Jiwu Shu, Max: A Multicore-Accelerated File System for Flash Storage, The 2021 *USENIX Annual Technical Conference (USENIX ATC)*, JULY 14–16, 2021 (a virtual event), (Acceptance rate: 64/341=18.77%)
- [5] Shiyao Lin, Guowen Gong, Zhirong Shen, Patrick P. C. Lee, Jiwu Shu, Boosting Full-Node Repair in Erasure-Coded Storage, The 2021 *USENIX Annual Technical Conference (USENIX ATC)*, JULY 14–16, 2021 (a virtual event), (Acceptance rate: 64/341=18.77%)
- [6] Fan Yang, Youmin Chen, Youyou Lu, Qing Wang, Jiwu Shu, Aria: Tolerating Skewed Workloads in Secure In-memory Key-value Stores, The 37th *IEEE International Conference on Data Engineering (ICDE)*, 19-22 April 2021
- [7] Youmin Chen, Youyou Lu, Bohong Zhu, Andrea Arpaci-Dusseau, Remzi Arpaci-Dusseau, Jiwu Shu, Scalable Persistent Memory File System with Kernel-User Space Collaboration, 19th *USENIX Conference on File and Storage Technologies (FAST)*, February 23-25, 2021, Santa Clara, California, USA. (Acceptance Rate: 28/130=21.54%)
- [8] Qing Wang, Youyou Lu, Erci Xu, Junyu Li, Youmin Chen, Jiwu Shu, Concordia: Distributed Share memory with In-Network Cache Coherence, 19th *USENIX Conference on File and Storage Technologies (FAST)*, February 22-25, 2021, Santa Clara, California, USA. (Acceptance Rate: 28/130=21.54%)
- [9] Xiaojian Liao, Youyou Lu, Erci Xu, Jiwu Shu, Write Dependency Disentanglement with HORAE, the 14th *USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, November 4–6, 2020. Pp.549-565 (Acceptance rate: 70/398=17.59%)
- [10] Minhui Xie, Kai Ren, Youyou Lu, JiaZhen Lin, Guangxu Yang, Bihai Wu, Fan Yang, Wanhong Xu, Jiwu Shu, Kraken: Memory Efficient Continual Learning for Large-Scale Real-Time

- Recommendation, *The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC)*, Nov. 16-19, 2020, Denver, Co
- [11] Youmin Chen, Youyou Lu, Bohong Zhu, Jiwu Shu, Tree: a Persistent B+-Tree with Low Tail Latency, *Proceedings of Very Large Data Bases (VLDB)*, Aug 31, 2020 - Sep 4, 2020, Tokyo, Japan
- [12] Qing Wan, Youyou Lu, Zhongjie Wu, Fan Yang, Jiwu Shu, Improving the Concurrency Performance of Persistent Memory Transactions on Multicores, the 56th Annual *Design Automation Conference (DAC)*, July 19, 2020, San Francisco, California, USA (Acceptance rate: 228/984=23.17%)
- [13] Yang Zhe, Youyou Lu, Erci Xu, Jiwu Shu, CoinPurse: A Device-Assisted File System with Dual Interfaces, the 56th Annual *Design Automation Conference (DAC)*, July 19, 2020, San Francisco, California, USA. (Acceptance rate: 228/984=23.17%)
- [14] Youmin Chen, Youyou Lu, Fan Yang, Qing Wang, Yang Wang, Jiwu Shu, FlatStore: an Efficient Log-Structured Key-Value Storage Engine for Persistent Memory, The 25th *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, March 16-20, 2020, Lausanne, Switzerland (Acceptance rate: 86/476=18.07%)
- [15] Fan Yang, Youyou Lu, Youmin Chen, Haiyu Mao, Jiwu Shu, No Compromises: Secure NVM with Crash Consistency, Write-Efficiency and High-Performance, the 55th Annual *Design Automation Conference (DAC)*, June 2-6, 2019, Las Vegas, NV, USA (Acceptance rate: 202/815=24.79%)
- [16] Fei Li, Youyou Lu, Zhongjie Wu, Jiwu Shu, ASCache: An Approximate SSD Cache for Error-Tolerant Applications, the 55th Annual *Design Automation Conference (DAC)*, June 2-6, 2019, Las Vegas, NV, USA (Acceptance rate: 202/815=24.79%)
- [17] Youmin Chen, Youyou Lu, Jiwu Shu, Scalable RDMA RPC on Reliable Connection with Efficient Resource Sharing, *European Conference on Computer Systems (EuroSys)*, Dresden, Germany, Mar 25, 2019 - Mar 28, 2019 (Acceptance rate: 45/207=21.74%)
- [18] Haiyu Mao, Mingcong Song, Tao Li, Yutiing Dai, Jiwu Shu, LerGAN: A Zero-free, Low Data Movement and PIM-based GAN Architecture, The 51st Annual IEEE/ACM *International Symposium on Microarchitecture (MICRO)*, Fukuoka City, Japan, 2018, Oct 20 – 24, pp.670-682 (acceptance rate: 74/351=21.06%)
- [19] Siyang Li, Youyou Lu, Jiwu Shu, Fenlin Liu, Yang Hu, Tao Li, LocoFS: A Loosely-Coupled Metadata Service for Distributed File System, The *International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, Nov. 12-17, 2017, Denver, Co (Acceptance Rate: 61/327=18.65%)
- [20] Youyou Lu, Jiwu Shu, Youmin Chen, Tao Li, Octopus: an RDMA-enabled Distribute Persistent Memory File System, 2017 *USENIX Annual Technical Conference (USENIX ATC)*, July 12–14, 2017, Santa Clara, CA, USA, pp.773-785 (acceptance rate: 60/283=21%)
- [21] Qingda Hu, Jinglei Ren, Anirudh Badam, Jiwu Shu, Tomas Moscibroda, Log-Structured Non-Volatile Main Memory, 2017 *USENIX Annual Technical Conference (USENIX ATC)*, July 12–14, 2017, Santa Clara, CA, USA, pp.703-717 (acceptance rate: 60/283=21%)
- [22] Jiacheng Zhang, Jiwu Shu, Youyou Lu, ParaFS: A Log-Structured File System to Exploit the Internal Parallelism of Flash Devices, 2016 *USENIX Annual Technical Conference (USENIX*

- ATC), June 22–24, 2016, Denver, Colorado, USA (Acceptance Rate: 19%)
- [23] Zhirong Shen, Jiwu Shu, Patrick P.C. Lee, Reconsidering Single Failure Recovery in Cluster File Systems, *The 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 2016, June 28-July 1st, Toulouse, France (Acceptance Rate:  $58/259=22.39\%$ )
- [24] Zewei Li, Yongpan Liu, Daming Zhang, Chun Jason Xue, Zhangyuan Wang, Xin Shi, Wenyu Sun, Jiwu Shu, Huazhong Yang, HW/SW Co-design of Nonvolatile IO System in Energy Harvesting Sensor Nodes for Optimal Data Acquisition, *Design Automation Conference (DAC)*, 5–9 June 2016, Austin, USA (Acceptance Rate:  $152/876=17.4\%$ )
- [25] Jiaxin Ou, Jiwu Shu, Youyou Lu, A High Performance File System for Non-Volatile Main Memory, *The European Conference on Computer Systems (EuroSys)*, 18th-21st, April, 2016, London, UK (Acceptance Rate:  $38/180=21.11\%$ )
- [26] Zhirong Sheng, Jiwu Shu, Patrick Lee, Seek-Efficient I/O Optimization in Single Failure Recovery for XOR-Coded Storage Systems, *The 34th International Symposium on Reliable Distributed Systems (SRDS)*, September 28-October 1, 2015, Quebec Canada (**Nominated Best Papers**).
- [27] Yongpan Liu, Zewei Li, Hehe Li, Yiqun Wang, Xueqing Li, Kaisheng Ma, Shuangchen Li, Meng-Fan Chang, Sampson John, Yuan Xie, Jiwu Shu, Huazhong Yang, *Design Automation Conference (DAC)*, Moscone Center, San Francisco, CA, June 7-11, 2015, Article No.150, pp.1-6
- [28] Chao Zhang, Guangyu Sun, Xian Zhang, Weiqi Zhang, Weisheng Zhao, Tao Wang, Yun Liang, Yongpan Liu, Yu Wang, Jiwu Shu, Hi-fi Playback: Tolerating Position Errors in Shift Operations of Racetrack Memory, *The 42nd International Symposium on Computer Architecture (ISCA)*, Portland, OR, USA, June 13-17, 2015, pp. 694-706. (Acceptance Rate:  $58/305=19.02\%$ )
- [29] Louyou Lu, Jiwu Shu, Long Sun, Blurred Persistence in Transactional Persistent Memory, *The 31st International Conference on Massive Storage Systems and Technologies (MSST)*, June 1-5, 2015, Santa Clara, California, USA (**Nominated Best Papers**)
- [30] Fan Jie, Jiang Song, Shu Jiwu, Sun Long, Hu Qingda, WL-Reviver: A Framework for Reviving any Wear-Leveling Techniques in the Face of Failures on Phase Change Memory, *The Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, June 23 - 26, 2014, Atlanta, Georgia USA, pp.228-239. (Acceptance Rate:  $49/181 = 27.07\%$ )
- [31] Zhirong Sheng, Jiwu Shu, HV Code: An MDS Code to Improve Efficiency and Reliability of RAID-6 Systems, *The Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, June 23 - 26, 2014, Atlanta, Georgia USA, pp.550-56. (Acceptance Rate:  $49/181 = 27.07\%$ )
- [32] Lu Youyou, Shu Jiwu, Wang Wei, ReconFS: A Reconstructable File System on Flash Storage, *the 12th USENIX Conference on File and Storage Technologies (FAST)*, February 12-15, 2014, San Jose, CA, USA, pp. 75-88. (Acceptance Rate:  $21/133=15.79\%$ )
- [25] Jie Fan, Song Jiang, Jiwu Shu, Youhui Zhang, Weimin Zheng, Partitioning Data Block for Efficient Recovery of Stuck-At-Faults in Phase Change Memory, *The 46th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, December 7-11, 2013, Davis, California, pp. 433-444. (Acceptance Rate:  $39/239=16.32\%$ )

- [33] Youyou Lu, Jiwu Shu, Weimin Zheng, Extending the Lifetime of Flash-based Storage through Reducing Write Amplification from File Systems, 11th *USENIX Conference on File and Storage Technologies (FAST)*, February 12-15, 2013, San Jose, CA, USA, pp. 257-270. (Acceptance Rate: 20/127=15.75%)

## ● 代表性期刊论文

- [1] Haiyu Mao, Jiwu Shu, Mingcong Song, Tao Li, LrGAN: A Compact and Energy Efficient PIM-based Architecture for GAN Training, *IEEE Transactions on Computers (TC)*, Volume: 70, Issue: 9, SEPTEMBER 2021: 1427-1442
- [2] Jiwu Shu, Youmin Chen, Qing Wang, Bohong Zhu, Junru Li, Youyou Lu, TH-DPMS: Design and Implementation of an RDMA-enabled Distributed Persistent Memory Storage System, *ACM Transactions on Storage (TOS)*, 16, 4 Article 24 (November 2020), 31 pages.
- [3] Fan Yang, Youmin Chen, Haiyu Mao, Youyou, Jiwu Shu, ShieldNVM: An Efficient and Fast Recoverable System for Secure Non-Volatile Memory, *ACM Transactions on Storage (TOS)*, Vol.16, No.2, Article 12, Publication date: May 2020.
- [4] Jiwu Shu, Fei Li, Siyang Li, Youyou Lu, Towards Unaligned Writes Optimization in Cloud Storage with High-performance SSDs, *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, Vol. 31, No.12 December 2020, pp.2923-2937.
- [5] Zhirong Shen, Patrick P.C. Lee, Jiwu Shu, and Wenzhong Guo, Cross-Rack-Aware Single Failure Recovery for Clustered File Systems, *IEEE Transactions on Dependable and Secure Computing (TDSC)*, Vol.17, Issue2, pp.248-261, MARCH/APRIL 2020.
- [6] Youyou Lu, Jiwu Shu, Jiacheng Zhang, Mitigating Synchronous I/O Overhead in File Systems on Open-Channel SSDs, *ACM Transactions on Storage (TOS)*, 15(3), Article 17 :1-25, May 2019
- [7] Youmin Chen, Youyou Lu, Pei Chen, Jiwu Shu, Efficient and Consistent NVMM Cache for SSD-based File System, *IEEE Transactions on Computers (TC)*, 68(8): 1147-1158, August, 2019.
- [8] Zhirong Shen, Patrick P. C. Lee, Jiwu Shu and Wenzhong Guo, Correlation-Aware Stripe Organization for Efficient Writes in Erasure-Coded Storage: Algorithms and Evaluation, *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 30(7):1552-1564, July 2019.
- [9] Siyang Li, Fenlin Liu, Jiwu Shu, Youyou Lu, Tao Li, Yang Hu, A Flattened Metadata Service for Distributed File Systems, *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 29(12):2641-2657, December 2018.
- [10] Zhirong Shen, Patrick P. C. Lee, Jiwu Shu and Wenzhong Guo, Encoding-Aware Data Placement for Efficient Degraded Reads in XOR-Coded Storage Systems: Algorithms and Evaluation, *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 29(12):2757-2770, December 2018.
- [11] Youmin Chen, Jiwu Shu, Jiaxin Ou, Youyou Lu, HiNFS: A Persistent Memory File System with both Buffering and Direct-Access, *ACM Transactions on Storage (TOS)*, Vol.14, Number 1, April 2018, 4:1-4:30.
- [12] Zhirong Shen, Jiwu Shu, Patrick P.C. Lee, and Yingxun Fu, Seek-Efficient I/O Optimization in Single Failure Recovery for XOR-Coded Storage Systems, *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 28(3):877-890, 2017.

- [13] Yingxun Fu, Jiwu Shu, Xianghong Luo, Zhirong Shen, Qingda Hu, Short Code: An Efficient RAID-6 MDS Code for Optimizing Degraded Reads and Partial Stripe Writes, *IEEE Transactions on Computers (TC)*, 66(1), pp. 127-137, 2017.
- [14] Jiacheng Zhang, Youyou Lu, Jiwu Shu, Xiongjun Qin, FlashKV: accelerating KV performance with open-channel SSDs, *ACM Transactions on Embedded Computing Systems (ACM TECS)*, Vol. 16, No. 5s, Article 139:1-19 (September 2017)
- [15] Youyou Lu, Jiwu Shu, Jia Guo, Peng Zhu, Supporting System Consistency with Differential Transactions in Flash-based SSDs, *IEEE Transactions on Computers (TC)*, 65(2), pp. 627-639, 2016.
- [16] Louyou Lu, Jiwu Shu, Long Sun, Blurred Persistence: Efficient Transactions in Persistent Memory, *ACM Transactions on Storage (TOS)*, Vol.12, No. 1, Article 3, pp. 1-29, January 2016.
- [17] Zhirong Shen, Jiwu Shu and Yingxun Fu, Parity-switched Data Placement: Optimizing Partial Stripe Writes in XOR-Coded Storage Systems, *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 27(11), pp. 3311-3322, 2016.
- [18] Zhirong Shen, Jiwu Shu, HV Code: An All-Around MDS Code for RAID-6 Storage Systems, *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 27(6), pp.1674-1686, JUNE 2016.
- [19] Yingxun Fu, Jiwu Shu, Zhang Guangyan, Reconsidering Single Disk Failure Recovery for Erasure Coded Storage Systems: Optimizing The Load Balancing in Stack-Level, *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 27(5):1457-1469, 2016.
- [20] Guangyan Zhang, Guiyong Wu, shupeng Wang, Jiwu Shu, Weimin Zheng, Keqin Li, CaCo: An Efficient Cauchy Coding Approach for Cloud Storage Systems, *IEEE Transactions on Computers (TC)*, 65(2), pp.435-447, 2016.
- [21] Youyou Lu, Jiwu Shu, Jia Guo, Shuai Li, and Onur Mutlu, High-Performance and Lightweight Transaction Support in Flash-based SSDs, *IEEE Transactions on Computers (TC)*, 64(10), pp.2819-2832, 2015.
- [22] Zhang Guangyan, Wang Jigang, Li Keqin, Shu Jiwu, and Zheng Weimin, Redistribute Data to Regain Load Balance during RAID-4 Scaling, *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, Vol.25, No.7, pp.1-11, JULY 2014.
- [23] Yi Letian, Shu Jiwu, Zhao Ying, Qian Yingjin, Lu Youyou, Zheng Weiming, Design and implementation of an asymmetric block-based parallel file system, *IEEE Transactions on Computers (TC)*, Vol.62, No.7, pp.1723-1735, JULY 2014.
- [24] Xianghong Luo, Jiwu Shu, Generalized X-code: An efficient RAID-6 code for arbitrary size of disk array, *ACM Transactions on Storage (TOS)*, Vol.8, No.3, Article 10, pp.1-19, 2012.
- [25] Wang Yang, Shu Ji-Wu, Zhang Guang-Yan, Xue Wei, SOPA: Selecting the Optimal Policy Adaptively for a cache system, *ACM Transactions on Storage (TOS)*, Vol. 6, No. 2, Article 7, pp. 1-18, 2010.
- [26] Guangyan Zhang, Weimin Zheng, and Jiwu Shu. ALV: A New Data Redistribution Approach to RAID-5 Scaling, *IEEE Transactions on Computers (TC)*, Vol. 59(3), pp. 345-357, 2010.
- [27] Guangyan Zhang, Jiwu Shu, Wei Xue, and Weimin Zheng, SLAS: An Efficient Approach to Scaling Round-robin Striped Volumes, *ACM Transactions on Storage (TOS)*, 2007, 3(1), Article3, pp.1-39.



- [28] Guangyan Zhang, Jiwu Shu, Wei Xue, and Weimin Zheng, Design and Implementation of an Out-of-Band Virtualization System for Large SANs, *IEEE Transactions on Computers (TC)*, 2007, 56(12), pp.1654-1665.
- [29] Shu Ji-wu, Li Bigang, Zheng Wei-min, Design and Implementation of a SAN System Based on the Fiber Channel Protocol, *IEEE Transactions on Computers (TC)*, 54(4), 2005, pp. 439-448.
- [30] JIWU SHU, WEI XUE, WEIMING ZHENG, A Parallel Transient Stability Simulation for Power System, *IEEE Transactions on Power Systems (TOPS)*, 2005, 20(4), pp.1709-1717.
- [31] Li Mingqiang, Shu Jiwu, Preventing Silent Data Corruptions from Propagating During Data Reconstruction, *IEEE Transactions on Computers (TC)*, 59(12), pp.1611-1624, 2010 (Featured Article).
- [32] Li Mingqiang, Shu Jiwu, DACO: A High Performance Disk Architecture Designed Specially for Large Scale Erasure Coded Storage Systems, *IEEE Transactions on Computers (TC)*, 59(10), pp.1350-1362, 2010.
- [33] Li Mingqiang, Shu Jiwu, Zheng Weimin, GRID Codes: Strip-based Erasure Codes with High Fault Tolerance for Storage Systems, *ACM Transactions on Storage (TOS)*, 2009, 4(4), Article15, pp.1-22.