

Curriculum Vitae of Patrick P. C. Lee

Department of Computer Science and Engineering
The Chinese University of Hong Kong

January 26, 2024

1 Background Information

1.1 Personal Data

Name: Patrick P. C. Lee (Patrick Pak-Ching Lee)
Address: Room 108, Ho Sin Hang Engineering Building
Department of Computer Science and Engineering
The Chinese University of Hong Kong
Shatin, Hong Kong
Tel: (852) 3943 4260
Fax: (852) 2603 5024
Email: pclee@cse.cuhk.edu.hk
URL: <http://www.cse.cuhk.edu.hk/~pclee>
ORCID: 0000-0002-4501-4364 (<https://orcid.org/0000-0002-4501-4364>)

1.2 Education

- **Columbia University, New York, NY** September 2003 - August 2008
Ph.D. in Computer Science
Thesis: "Achieving Network Robustness with Diversity Overlay Routing"
Advisors: Prof. Vishal Misra and Prof. Dan Rubenstein
- **The Chinese University of Hong Kong, Hong Kong** August 2001 - July 2003
M.Phil. in Computer Science and Engineering
Thesis: "Distributed and Collaborative Key Agreement Protocols with Authentication and Implementation for Dynamic Peer Groups"
Advisor: Prof. John C. S. Lui
- **The Chinese University of Hong Kong, Hong Kong** September 1997 - July 2001
Bachelor of Engineering (First Class Honors), Information Engineering
- **The University of British Columbia, Canada** September 1999 - May 2000
Visiting Student, Electrical and Computer Engineering

1.3 Experience

- **The Chinese University of Hong Kong, Hong Kong**
Professor, Department of Computer Science and Engineering July 2021 - present
Vice Chairman (Undergraduate), Department of Computer Science and Engineering August 2022 - present
Associate Professor, Department of Computer Science and Engineering July 2015 - July 2021
Assistant Professor, Department of Computer Science and Engineering July 2009 - July 2015
- **Alcatel-Lucent, Murray Hill, NJ**
Consultant April 2007 - June 2011
- **University of Massachusetts, Amherst, MA**
Postdoctoral researcher, Department of Computer Science September 2008 - July 2009
- **Bell Laboratories, Lucent Technologies, Murray Hill, NJ**
Summer intern June 2006 - August 2006
- **Bell Laboratories, Lucent Technologies, Holmdel, NJ**
Summer intern June 2005 - August 2005

2 Contributions in Research

My general research interests are in various applied/systems topics on improving the dependability of storage systems and computer networks. By dependability, I refer to the fault tolerance, recovery, security, and performance guarantees that need to be achieved in order to maintain the correctness and performance of regular operations. My research approach emphasizes system design and implementation, backed by extensive empirical experiments and possibly theoretical analysis.

My current research blends storage and networking, with specific emphasis on two major themes: (i) **dependability of storage systems**, in which we explore practical data management mechanisms that can address the dependability of different types of storage architectures, including clouds, data centers, disk arrays, SSDs, and memory, in the face of failures and attacks; and (ii) **dependability of computer networks**, in which we explore new data structures and machine learning techniques, so as to achieve accurate network-wide measurement at scale as well as real-time and robust data analytics in distributed stream processing.

2.1 Publications

Google Scholar Citations: <https://scholar.google.com.hk/citations?user=gyRtVVEAAAAJ>

Note: ¹ = CUHK students under my supervision, ² = CUHK postdocs, ³ = visiting students. The work was done when they were affiliated with CUHK.

2.1.1 Conference/Workshop/Poster Publications

1. Qingxiu Liu¹, Qun Huang, Xiang Chen, Sa Wang, Wenhao Wang, Shujie Han, and Patrick P. C. Lee. “PP-Stream: Toward High-Performance Privacy-Preserving Neural Network Inference via Distributed Stream Processing.” Proceedings of the 40th IEEE International Conference on Data Engineering (**ICDE 2024**), Utrecht, Netherlands, May 2024.
2. Yanjing Ren¹, Yuanming Ren¹, Xiaolu Li, Yuchong Hu, Jingwei Li, and Patrick P. C. Lee. “ELECT: Enabling Erasure Coding Tiering for LSM-tree-based Storage.” Proceedings of the 22nd USENIX Conference on File and Storage Technologies (**FAST 2024**), Santa Clara, CA, USA, February 2024. (AR: 22/123 = 17.9%)
3. Qiuping Wang¹ and Patrick P. C. Lee. “ZapRAID: Toward High-Performance RAID for ZNS SSDs via Zone Append.” Proceedings of the 14th ACM SIGOPS Asia-Pacific Workshop on Systems (**APSys 2023**), Seoul, South Korea, August 2023. (AR: 13/32 = 40.6%)
4. Mi Zhang, Qihan Kang, and Patrick P. C. Lee. “Minimizing Network and Storage Costs for Consensus with Flexible Erasure Coding.” Proceedings of the 52nd International Conference on Parallel Processing (**ICPP 2023**), Salt Lake City, UT, USA, August 2023. (AR: 82/278 = 29.5%)
5. Siyuan Sheng¹, Huancheng Puyang¹, Qun Huang, Lu Tang, Patrick P. C. Lee. “FarReach: Write-back Caching in Programmable Switches.” Proceedings of the 2023 USENIX Annual Technical Conference (**USENIX ATC 2023**), Boston, USA, July 2023. (AR: 65/353 = 18.4%)
6. Kaicheng Tang¹, Keyun Cheng¹, Helen H. W. Chan², Xiaolu Li, Patrick P. C. Lee, Yuchong Hu, Jie Li, and Ting-Yi Wu. “Balancing Repair Bandwidth and Sub-packetization in Erasure-Coded Storage via Elastic Transformation.” Proceedings of IEEE International Conference on Computer Communications (**INFOCOM 2023**), New York, USA, May 2023. (AR: 252/1312 = 19.2%)
7. Jingwei Li, Yanjing Ren¹, Patrick P. C. Lee, Yuyu Wang, Ting Chen, Xiaosong Zhang. “FeatureSpy: Detecting Learning-Content Attacks via Feature Inspection in Secure Deduplicated Storage.” Proceedings of IEEE International Conference on Computer Communications (**INFOCOM 2023**), New York, USA, May 2023. (AR: 252/1312 = 19.2%)
8. Xiaolu Li, Keyun Cheng¹, Kaicheng Tang¹, Patrick P. C. Lee, Yuchong Hu, Dan Feng, Jie Li, and Ting-Yi Wu. “ParaRC: Embracing Sub-Packetization for Repair Parallelization in MSR-Coded Storage.” Proceedings of the 21st USENIX Conference on File and Storage Technologies (**FAST 2023**), Santa Clara, CA, USA, February 2023. (AR: 28/122 = 23.0%)

9. Zhinan Cheng¹, Shujie Han, Patrick P. C. Lee, Xin Li, Jiongzhou Liu, and Zhan Li. “An In-Depth Correlative Study Between DRAM Errors and Server Failures in Production Data Centers.” Proceedings of the 41st International Symposium on Reliable Distributed Systems (**SRDS 2022**), Vienna, Austria, September 2022. (AR: 24/105 = 22.9%)
10. Zuoru Yang¹, Jingwei Li, and Patrick P. C. Lee. “Secure and Lightweight Deduplicated Storage via Shielded Deduplication-Before-Encryption.” Proceedings of the 2022 USENIX Annual Technical Conference (**USENIX ATC 2022**), July 2022. (AR: 64/393 = 16.3%)
11. Qiaori Yao, Yuchong Hu, Xinyuan Tu, Patrick P. C. Lee, Dan Feng, Xia Zhu, Xiaoyang Zhang, Zhen Yao, and Wenjia Wei. “PivotRepair: Fast Pipelined Repair for Erasure-Coded Hot Storage.” Proceedings of the 42nd IEEE International Conference on Distributed Computing Systems (**ICDCS 2022**), July 2022. (AR: 113/573 = 19.7%)
12. Hao Shi, Hanxu Hou, Yunghsiang S. Han, Patrick P. C. Lee, Zhengyi Jiang, Zhongyi Huang, Bo Bai. “New Piggybacking Codes with Lower Repair Bandwidth for Any Single-Node Failure.” Proceedings of the 2022 IEEE International Symposium on Information Theory (**ISIT 2022**), June 2022.
13. Si Wu, Qingpeng Du, Patrick P. C. Lee, Yongkun Li, Yinlong Xu. “Optimal Data Placement for Stripe Merging in Locally Repairable Codes.” Proceedings of IEEE International Conference on Computer Communications (**INFOCOM 2022**), May 2022. (AR: 225/1129 = 19.9%)
14. Jingwei Li, Guoli Wei, Jiacheng Liang, Yanjing Ren, Patrick P. C. Lee, Xiaosong Zhang. “Revisiting Frequency Analysis against Encrypted Deduplication via Statistical Distribution.” Proceedings of IEEE International Conference on Computer Communications (**INFOCOM 2022**), May 2022. (AR: 225/1129 = 19.9%)
15. Qiang Zhang, Yongkun Li, Patrick P. C. Lee, Yinlong Xu, and Si Wu. “DEPART: Replica Decoupling for Distributed Key-Value Storage.” Proceedings of the 20th USENIX Conference on File and Storage Technologies (**FAST 2022**), February 2022. (AR: 28/130 = 21.5%)
16. Qiuping Wang¹, Jinhong Li¹, Patrick P. C. Lee, Tao Ouyang, Chao Shi, Lilong Huang. “Separating Data via Block Invalidation Time Inference for Write Amplification Reduction in Log-Structured Storage.” Proceedings of the 20th USENIX Conference on File and Storage Technologies (**FAST 2022**), February 2022. (AR: 28/130 = 21.5%)
17. Siyuan Sheng¹, Qun Huang, and Patrick P. C. Lee. “DeltaINT: Toward General In-band Network Telemetry with Extremely Low Bandwidth Overhead.” Proceedings of the 29th IEEE International Conference on Network Protocols (**ICNP 2021**), November 2021. (AR: 38/154 = 24.7%)
18. Hanxu Hou and Patrick P. C. Lee. “STAR+ Codes: Triple-Fault-Tolerant Codes with Asymptotically Optimal Updates and Efficient Encoding/Decoding.” Proceedings of the 2021 IEEE Information Theory Workshop (**ITW 2021**), October 2021.
19. Zhinan Cheng¹, Lu Tang, Qun Huang, and Patrick P. C. Lee. “Enabling Low-Redundancy Proactive Fault Tolerance for Stream Machine Learning via Erasure Coding.” Proceedings of the 40th International Symposium on Reliable Distributed Systems (**SRDS 2021**), September 2021. (AR: 27/106 = 25.5%)
20. Rui Li, Zhinan Cheng², Patrick P. C. Lee, Pinghui Wang, Yi Qiang, Lin Lan, Cheng He, Jinlong Lu, Mian Wang, Xinquan Ding. “Automated Intelligent Healing in Cloud-Scale Data Centers.” Proceedings of the 40th International Symposium on Reliable Distributed Systems (**SRDS 2021**), September 2021. (AR: 27/106 = 25.5%)
21. Yanjing Ren, Jingwei Li, Zuoru Yang¹, Patrick P. C. Lee, and Xiaosong Zhang. “Accelerating Encrypted Deduplication via SGX.” Proceedings of the 2021 USENIX Annual Technical Conference (**USENIX ATC 2021**), July 2021. (AR: 64/341 = 18.8%)
22. Yongkun Li, Zhen Liu, Patrick P. C. Lee, Jiayu Wu, Yinlong Xu, Yi Wu, Liu Tang, Qi Liu, and Qiu Cui. “Differentiated Key-Value Storage Management for Balanced I/O Performance.” Proceedings of the 2021 USENIX Annual Technical Conference (**USENIX ATC 2021**), July 2021. (AR: 64/341 = 18.8%)

23. Shiyao Lin, Guowen Gong, Zhirong Shen, Patrick P. C. Lee, and Jiwu Shu. “Boosting Full-Node Repair in Erasure-Coded Storage.” Proceedings of the 2021 USENIX Annual Technical Conference (**USENIX ATC 2021**), July 2021. (AR: 64/341 = 18.8%)
24. Hanxu Hou and Patrick P. C. Lee. “Generalized Rack-aware Regenerating Codes for Jointly Optimal Node and Rack Repairs.” Proceedings of the 2021 IEEE International Symposium on Information Theory (**ISIT 2021**), July 2021.
25. Qiaori Yao, Yuchong Hu, Liangfeng Cheng, Patrick P. C. Lee, Dan Feng, Weichun Wang, Wei Chen. “StripeMerge: Efficient Wide-Stripe Generation for Large-Scale Erasure-Coded Storage.” Proceedings of the 41st IEEE International Conference on Distributed Computing Systems (**ICDCS 2021**), July 2021. (AR: 97/498 = 19.8%)
26. Fan Xu, Shujie Han¹, Patrick P. C. Lee, Yi Liu, Cheng He, and Jiongzhou Liu. “General Feature Selection for Failure Prediction in Large-scale SSD Deployment.” Proceedings of the 51st IEEE/IFIP International Conference on Dependable Systems and Networks (**DSN 2021**), June 2021. (AR: 48/295 = 16.3%)
27. Guowen Gong, Zhirong Shen, Suzhen Wu, Xiaolu Li², and Patrick P. C. Lee. “Optimal Rack-Coordinated Updates in Erasure-Coded Data Centers.” Proceedings of IEEE International Conference on Computer Communications (**INFOCOM 2021**), May 2021. (AR: 252/1266 = 19.9%)
28. Shujie Han¹, Patrick P. C. Lee, Fan Xu, Yi Liu, Cheng He, and Jiongzhou Liu. “An In-Depth Study of Correlated Failures in Production SSD-Based Data Centers.” Proceedings of the 19th USENIX Conference on File and Storage Technologies (**FAST 2021**), February 2021. (AR: 28/130 = 21.5%)
29. Yuchong Hu, Liangfeng Cheng, Qiaori Yao, Patrick P. C. Lee, Weichun Wang, and Wei Chen. “Exploiting Combined Locality for Wide-Stripe Erasure Coding in Distributed Storage.” Proceedings of the 19th USENIX Conference on File and Storage Technologies (**FAST 2021**), February 2021. (AR: 28/130 = 21.5%)
30. You Wu, Hanxu Hou, Yunghsiang S. Han, Patrick P. C. Lee, and Guojun Han. “Generalized Expanded-Blaum-Roth Codes and Their Efficient Encoding/Decoding.” Proceedings of the 2020 IEEE Global Communications Conference (**GLOBECOM 2020**), Taipei, Taiwan, December 2020.
31. Shujie Han¹, Patrick P. C. Lee, Zhirong Shen, Cheng He, Yi Liu, and Tao Huang. “Toward Adaptive Disk Failure Prediction via Stream Mining.” Proceedings of the 40th IEEE International Conference on Distributed Computing Systems (**ICDCS 2020**), Singapore, November 2020. (AR: 105/584 = 18.5%)
32. Jinhong Li¹, Qiuping Wang¹, Patrick P. C. Lee, and Chao Shi. “An In-Depth Analysis of Cloud Block Storage Workloads in Large-Scale Production.” Proceedings of 2020 IEEE International Symposium on Workload Characterization (**IISWC 2020**), Beijing, China, October 2020. (AR: 26/70 = 37.1%)
33. Si Wu², Zhirong Shen, and Patrick P. C. Lee. “Enabling I/O-Efficient Redundancy Transitioning in Erasure-Coded KV Stores via Elastic Reed-Solomon Codes.” Proceedings of the 39th International Symposium on Reliable Distributed Systems (**SRDS 2020**), Shanghai, China, September 2020. (AR: 33/131 = 25.2%)
34. Qun Huang, Haifeng Sun, Patrick P. C. Lee, Wei Bai, Feng Zhu, and Yungang Bao. “OmniMon: Re-architecting Network Telemetry with Resource Efficiency and Full Accuracy.” Proceedings of the ACM SIGCOMM 2020 conference (**SIGCOMM 2020**), New York, NY, USA, August 2020. (**Awarded with ACM artifact badges**)
35. Qiuping Wang¹, Jinhong Li¹, Wen Xia, Erik Kruus, Biplob Debnath, and Patrick P. C. Lee. “Austere Flash Caching with Deduplication and Compression.” Proceedings of the 2020 USENIX Annual Technical Conference (**USENIX ATC 2020**), Boston, MA, USA, July 2020. (AR: 65/348 = 18.6%)
36. Si Wu², Zhirong Shen², and Patrick P. C. Lee. “On the Optimal Repair-Scaling Trade-off in Locally Repairable Codes.” Proceedings of IEEE International Conference on Computer Communications (**INFOCOM 2020**), Toronto, Canada, July 2020. (AR: 268/1354 = 19.8%)
37. Lu Tang¹, Qun Huang, and Patrick P. C. Lee. “SpreadSketch: Toward Invertible and Network-Wide Detection of Superspreaders.” Proceedings of IEEE International Conference on Computer Communications (**INFOCOM 2020**), Toronto, Canada, July 2020. (AR: 268/1354 = 19.8%)

38. Hanxu Hou, Patrick P. C. Lee, and Yunghsiang S. Han. “Toward Optimality in Both Repair and Update via Generic MDS Code Transformation.” Proceedings of the 2020 IEEE International Symposium on Information Theory (**ISIT 2020**), Los Angeles, CA, USA, June 2020.
39. Hanxu Hou, Patrick P. C. Lee, and Yunghsiang S. Han. “Minimum Storage Rack-Aware Regenerating Codes with Exact Repair and Small Sub-Packetization.” Proceedings of the 2020 IEEE International Symposium on Information Theory (**ISIT 2020**), Los Angeles, CA, USA, June 2020.
40. Jingwei Li, Zuoru Yang¹, Yanjing Ren, Patrick P. C. Lee, and Xiaosong Zhang. “Balancing Storage Efficiency and Data Confidentiality with Tunable Encrypted Deduplication.” Proceedings of the 15th European Conference on Computer Systems (**EuroSys 2020**), Heraklion, Crete, Greece, April 2020. (AR: 43/234 = 18.4%)
41. Qiang Zhang³, Yongkun Li, Patrick P. C. Lee, Yinlong Xu, Qiu Cui, and Liu Tang. “UniKV: Toward High-Performance and Scalable KV Storage in Mixed Workloads via Unified Indexing.” Proceedings of the 36th IEEE International Conference on Data Engineering (**ICDE 2020**), Dallas, TX, USA, April 2020.
42. Liangfeng Cheng, Yuchong Hu, and Patrick P. C. Lee. “Coupling Decentralized Key-Value Stores with Erasure Coding.” Proceedings of the ACM Symposium on Cloud Computing 2019 (**SoCC 2019**), Santa Cruz, CA, USA, November 2019. (AR: 39/157 = 24.8%)
43. Hanxu Hou, Yunghsiang S. Han, Patrick P. C. Lee, and Qingfeng Zhou. “New Regenerating Codes over Binary Cyclic Codes.” Proceedings of the 2019 IEEE International Symposium on Information Theory (**ISIT 2019**), Paris, France, July 2019.
44. Hanxu Hou, Yunghsiang S. Han, and Patrick P. C. Lee. “Binary MDS Array Codes with Asymptotically Optimal Repair for All Columns.” Proceedings of the 28th International Conference on Computer Communications and Networks (**ICCCN 2019**) (Invited paper), Valencia, Spain, July 2019.
45. Zhirong Shen², Xiaolu Li¹, and Patrick P. C. Lee. “Fast Predictive Repair in Erasure-Coded Storage.” Proceedings of the 49th IEEE/IFIP International Conference on Dependable Systems and Networks (**DSN 2019**), Portland, Oregon, USA, June 2019. (AR: 54/212 = 21.4%)
46. Mi Zhang¹, Qiuping Wang¹, Zhirong Shen², and Patrick P. C. Lee. “Parity-Only Caching for Robust Straggler Tolerance.” Proceedings of 35th International Conference on Massive Storage Systems and Technology (**MSST 2019**) (Full paper), Santa Clara, CA, USA, May 2019. (AR: 21/71 = 29.6%)
47. Jingwei Li, Patrick P. C. Lee, Yanjing Ren, and Xiaosong Zhang. “Metadedup: Deduplicating Metadata in Encrypted Deduplication via Indirection.” Proceedings of 35th International Conference on Massive Storage Systems and Technology (**MSST 2019**) (Full paper), Santa Clara, CA, USA, May 2019. (AR: 21/71 = 29.6%)
48. Lu Tang¹, Qun Huang, and Patrick P. C. Lee. “MV-Sketch: A Fast and Compact Invertible Sketch for Heavy Flow Detection in Network Data Streams.” Proceedings of IEEE International Conference on Computer Communications (**INFOCOM 2019**), Paris, France, April-May 2019. (AR: 288/1464 = 19.7%)
49. Xiaolu Li¹, Runhui Li, Patrick P. C. Lee, and Yuchong Hu. “OpenEC: Toward Unified and Configurable Erasure Coding Management in Distributed Storage Systems.” Proceedings of the 17th USENIX Conference on File and Storage Technologies (**FAST 2019**), Boston, MA, USA, February 2019. (AR: 26/145 = 17.9%)
50. Qun Huang, Patrick P. C. Lee, and Yungang Bao. “SketchLearn: Relieving User Burdens in Approximate Measurement with Automated Statistical Inference.” Proceedings of the ACM SIGCOMM 2018 conference (**SIGCOMM 2018**), Budapest, Hungary, August 2018. (AR: 40/222 = 18.0%) (**Awarded with ACM artifact badges**)
51. Zhirong Shen² and Patrick P. C. Lee. “Cross-Rack-Aware Updates in Erasure-Coded Data Centers.” Proceedings of the 47th International Conference on Parallel Processing (**ICPP 2018**), Eugene, Oregon, USA, August 2018. (AR: 91/313 = 29.1%)
52. Helen H. W. Chan¹, Yongkun Li, Patrick P. C. Lee, and Yinlong Xu. “HashKV: Enabling Efficient Updates in KV Storage via Hashing.” Proceedings of the 2018 USENIX Annual Technical Conference (**USENIX ATC 2018**), Boston, MA, USA, July 2018. (AR: 76/378 = 20.1%)

53. Yuchong Hu, Xiaoyang Zhang, Patrick P. C. Lee, and Pan Zhou. “Generalized Optimal Storage Scaling via Network Coding.” Proceedings of IEEE International Symposium on Information Theory (**ISIT 2018**), Vali, Colorado, USA, June 2018.
54. Jun Wu³, Patrick P. C. Lee, Qi Li, Lujia Pan, and Jianfeng Zhang. “CellPAD: Detecting Performance Anomalies in Cellular Networks via Regression Analysis.” Proceedings of IFIP Networking (**Networking 2018**), Zurich, Switzerland, May 2018. (AR: 55/225 = 24.4%)
55. Zhiyong Wu, Yongkun Li, Patrick P. C. Lee, and Yinlong Xu. “Modeling SSD RAID Reliability under General Settings.” Proceedings of ACM International Conference on Computing Frontiers (**CF 2018**), Ischia, Italy, May 2018. (AR: 23/68 = 33.8%)
56. Xiaoyang Zhang, Yuchong Hu, Patrick P. C. Lee, Pan Zhou. “Toward Optimal Storage Scaling via Network Coding: From Theory to Practice.” Proceedings of IEEE International Conference on Computer Communications (**INFOCOM 2018**), Honolulu, HI, USA, April 2018. (AR: 309/1606 = 19.2%)
57. Mi Zhang¹, Shujie Han¹, and Patrick P. C. Lee. “A Simulation Analysis of Reliability in Erasure-Coded Data Centers.” Proceedings of the 36th IEEE International Symposium on Reliable Distributed Systems (**SRDS 2017**), Hong Kong, September 2017. (AR: 24/72 = 33.3%)
58. Zhirong Shen, Patrick P. C. Lee, Jiwu Shu, and Wenzhong Guo. “Correlation-Aware Stripe Organization for Efficient Writes in Erasure-Coded Storage Systems.” Proceedings of the 36th IEEE International Symposium on Reliable Distributed Systems (**SRDS 2017**), Hong Kong, September 2017. (AR: 24/72 = 33.3%)
59. Qun Huang, Xin Jin, Patrick P. C. Lee, Runhui Li, Lu Tang¹, Yi-Chao Chen, and Gong Zhang. “SketchVisor: Robust Network Measurement for Software Packet Processing.” Proceedings of the ACM SIGCOMM 2017 conference (**SIGCOMM 2017**), Los Angeles, CA, USA, August 2017. (AR: 36/250 = 14.4%)
60. Lujia Pan, Jianfeng Zhang, Patrick P. C. Lee, Hong Cheng, Cheng He, Caifeng He, and Keli Zhang. “An Intelligent Customer Care Assistant System for Large-Scale Cellular Network Diagnosis.” Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (**SIGKDD 2017**) (Applied data science track) (Poster presentation), Halifax, Canada, August 2017. (AR: (35+50)/396 = 21.5%)
61. Runhui Li, Xiaolu Li¹, Patrick P. C. Lee, and Qun Huang. “Repair Pipelining for Erasure-Coded Storage.” Proceedings of the 2017 USENIX Annual Technical Conference (**USENIX ATC 2017**), Santa Clara, CA, USA, July 2017. (AR: 60/283 = 21.2%)
62. Jingwei Li, Chuan Qin¹, Patrick P. C. Lee, and Xiaosong Zhang. “Information Leakage in Encrypted Deduplication via Frequency Analysis.” Proceedings of the 47th IEEE/IFIP International Conference on Dependable Systems and Networks (**DSN 2017**) (Regular paper), Denver, Colorado, USA, June 2017. (AR: 49/220 = 22.3%)
One of three nominations for the Best Paper Award.
63. Hanxu Hou², Patrick P. C. Lee, Yunghsiang S. Han, and Yuchong Hu. “Triple-Fault-Tolerant Binary MDS Array Codes with Asymptotically Optimal Repair.” Proceedings of IEEE International Symposium on Information Theory (**ISIT 2017**), Aachen, Germany, June 2017.
64. Matt M. T. Yiu¹, Helen H. W. Chan¹, and Patrick P. C. Lee. “Erasure Coding for Small Objects in In-Memory KV Storage.” Proceedings of the 10th ACM International Systems and Storage Conference (**SYSTOR 2017**), Haifa, Israel, May 2017. (AR: 17/45 = 37.8%)
65. Min Fu³, Patrick P. C. Lee, Dan Feng, Zuoning Chen, and Yu Xiao. “A Simulation Analysis of Reliability in Primary Storage Deduplication.” Proceedings of 2016 IEEE International Symposium on Workload Characterization (**IISWC 2016**), Providence, Rhode Island, USA, September 2016. (AR: 21/69 = 30.4%)
66. Zhirong Shen, Patrick P. C. Lee, Jiwu Shu, and Wenzhong Guo. “Encoding-Aware Data Placement for Efficient Degraded Reads in XOR-Coded Storage Systems.” Proceedings of the IEEE 35th Symposium on Reliable Distributed Systems (**SRDS 2016**), Budapest, Hungary, September 2016. (AR: 29/83 = 34.9%)

67. Yuchong Hu, Patrick P. C. Lee, and Xiaoyang Zhang. “Double Regenerating Codes for Hierarchical Data Centers.” Proceedings of IEEE International Symposium on Information Theory (**ISIT 2016**), Barcelona, Spain, July 2016.
68. Jingwei Li², Chuan Qin¹, Patrick P. C. Lee, and Jin Li. “Rekeying for Encrypted Deduplication Storage.” Proceedings of the 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (**DSN 2016**) (Regular paper), Toulouse, France, June 2016. (AR: 58/259 = 22.4%)
69. Yongkun Li, Helen H. W. Chan¹, Patrick P. C. Lee, and Yinlong Xu. “Elastic Parity Logging for SSD RAID Arrays.” Proceedings of the 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (**DSN 2016**) (Regular paper), Toulouse, France, June 2016. (AR: 58/259 = 22.4%)
70. Zhirong Shen, Jiwu Shu, and Patrick P. C. Lee. “Reconsidering Single Failure Recovery in Clustered File Systems.” Proceedings of the 46th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (**DSN 2016**) (Regular paper), Toulouse, France, June 2016. (AR: 58/259 = 22.4%)
71. Zhirong Shen, Patrick P. C. Lee, and Jiwu Shu. “Efficient Routing for Cooperative Data Regeneration in Heterogeneous Storage Networks.” Proceedings of the IEEE/ACM International Symposium on Quality of Service (**IWQoS 2016**) (Full paper), Beijing, China, June 2016. (AR: 27/131 = 20.6%)
72. Mingqiang Li², Chuan Qin¹, and Patrick P. C. Lee. “CDStore: Toward Reliable, Secure, and Cost-Efficient Cloud Storage via Convergent Dispersal.” Proceedings of the USENIX Annual Technical Conference (**USENIX ATC 2015**), Santa Clara, CA, USA, July 2015. (AR: 47/221 = 21.3%)
73. Runhui Li¹, Yuchong Hu², and Patrick P. C. Lee. “Enabling Efficient and Reliable Transition from Replication to Erasure Coding for Clustered File Systems.” Proceedings of the 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (**DSN 2015**) (Regular paper), Rio de Janeiro, Brazil, June 2015. (AR: 50/229 = 21.8%)
74. Min Xu¹, Yunfeng Zhu, Patrick P. C. Lee, and Yinlong Xu. “Even Data Placement for Load Balance in Reliable Distributed Deduplication Storage Systems.” Proceedings of the IEEE/ACM International Symposium on Quality of Service (**IWQoS 2015**) (Full paper), Portland, Oregon, USA, June 2015. (AR: 20/89 = 22.5%)
75. Qun Huang¹, Patrick P. C. Lee, Caifeng He, Jianfeng Qian, and Cheng He. “Fine-Grained Dissection of WeChat in Cellular Networks.” Proceedings of the IEEE/ACM International Symposium on Quality of Service (**IWQoS 2015**) (Full paper), Portland, Oregon, USA, June 2015. (AR: 20/89 = 22.5%)
76. Qun Huang¹, Patrick P. C. Lee, and Zhibin Zhang. “Exploiting Intra-Packet Dependency for Fine-Grained Protocol Format Inference.” Proceedings of IFIP Networking (**Networking 2015**), Toulouse, France, May 2015. (AR: 47/202 = 23.3%)
77. Runhui Li¹ and Patrick P. C. Lee. “Making MapReduce Scheduling Effective in Erasure-Coded Storage Clusters.” Proceedings of the 21st IEEE International Workshop on Local and Metropolitan Area Networks (**LAN-MAN 2015**) (Invited paper), Beijing, China, April 2015.
78. Yongkun Li, Patrick P. C. Lee, John C. S. Lui, and Yinlong Xu. “Impact of Data Locality on Garbage Collection in SSDs: A General Analytical Study.” Proceedings of the 6th ACM/SPEC International Conference on Performance Engineering (**ICPE 2015**) (Full paper), January 2015. (AR: 15/56 = 26.8%)
79. Runhui Li¹, Patrick P. C. Lee, and Yuchong Hu². “Degraded-First Scheduling for MapReduce in Erasure-Coded Storage Clusters.” Proceedings of the 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (**DSN 2014**) (Regular paper), Atlanta, Georgia, June 2014. (AR: 56/181 = 30.9%)
80. Mingqiang Li², Chuan Qin¹, Patrick P. C. Lee, and Jin Li. “Convergent Dispersal: Toward Storage-Efficient Security in a Cloud-of-Clouds.” Proceedings of the 6th USENIX Workshop on Hot Topics in Storage and File Systems (**HotStorage 2014**), Philadelphia, PA, June 2014. (AR: 15/57 = 26.3%)
81. Mingqiang Li² and Patrick P. C. Lee. “Toward I/O-Efficient Protection Against Silent Data Corruptions in RAID Arrays.” Proceedings of the 30th International Conference on Massive Storage Systems and Technology (**MSST 2014**) (Full paper), Santa Clara, CA, June 2014. (AR: 21/124 = 16.9%)

82. Qun Huang¹ and Patrick P. C. Lee. “LD-Sketch: A Distributed Sketching Design for Accurate and Scalable Anomaly Detection in Network Data Streams.” Proceedings of the 33rd Annual IEEE International Conference on Computer Communications (**INFOCOM 2014**), Toronto, Canada, April 2014. (AR: 320/1645 = 19.4%)
83. Zhuo Zhang, Zhibin Zhang, Patrick P. C. Lee, Yunjie Liu, and Gaogang Xie. “ProWord: An Unsupervised Approach to Protocol Feature Word Extraction.” Proceedings of the 33rd Annual IEEE International Conference on Computer Communications (**INFOCOM 2014**), Toronto, Canada, April 2014. (AR: 320/1645 = 19.4%)
84. Jeremy C. W. Chan¹, Qian Ding¹, Patrick P. C. Lee, and Helen H. W. Chan¹. “Parity Logging with Reserved Space: Towards Efficient Updates and Recovery in Erasure-coded Clustered Storage.” Proceedings of the 12th USENIX Conference on File and Storage Technologies (**FAST 2014**), Santa Calra, CA, February 2014. (AR: 24/133 = 18.0%)
85. Mingqiang Li² and Patrick P. C. Lee. “STAIR Codes: A General Family of Erasure Codes for Tolerating Device and Sector Failures in Practical Storage Systems.” Proceedings of the 12th USENIX Conference on File and Storage Technologies (**FAST 2014**), Santa Calra, CA, February 2014. (AR: 24/133 = 18.0%) **One of FAST 2014 papers selected for fast-track publication in ACM Transactions on Storage.**
86. Yongkun Li², Patrick P. C. Lee, and John C. S. Lui. “Stochastic Analysis on RAID Reliability for Solid-State Drives.” Proceedings of the 32nd IEEE International Symposium on Reliable Distributed Systems (**SRDS 2013**), Braga, Portugal, October 2013. (AR: 22/67 = 32.8%)
87. Chun-Ho Ng¹ and Patrick P. C. Lee “RevDedup: A Reverse Deduplication Storage System Optimized for Reads to Latest Backups.” The 4th ACM SIGOPS Asia-Pacific Workshop on Systems (**APSYS 2013**), Singapore, July 2013. (AR: 20/73 = 27.4%)
88. Yongkun Li², Patrick P. C. Lee, and John C. S. Lui. “Stochastic Modeling of Large-Scale Solid-State Storage Systems: Analysis, Design Tradeoffs and Optimization.” Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (**SIGMETRICS 2013**) (Full paper), Pittsburgh, PA, June 2013. (AR: 27/196 = 13.8%)
89. Runhui Li¹, Jian Lin¹, and Patrick P. C. Lee. “CORE: Augmenting Regenerating-Coding-Based Recovery for Single and Concurrent Failures in Distributed Storage Systems.” Proceedings of the 29th IEEE Conference on Massive Data Storage (**MSST 2013**) (Short paper), Long Beach, CA, May 2013. (AR: (14+15)/109 = 26.6%)
90. Yuchong Hu², Patrick P. C. Lee, and Kenneth W. Shum. “Analysis and Construction of Functional Regenerating Codes with Uncoded Repair for Distributed Storage Systems.” Proceedings of the 32nd Annual IEEE Conference on Computer Communications (**INFOCOM 2013**), Turin, Italy, April 2013. (AR: 280/1613 = 17%)
91. Shuang Zhao³, Patrick P. C. Lee, John C. S. Lui, Xiaohong Guan, Xiaobo Ma, and Jing Tao. “Cloud-based Push-Styled Mobile Botnets: A Case Study of Exploiting the Cloud to Device Messaging Service.” Proceedings of the 28th Annual Computer Security Applications Conference (**ACSAC 2012**), Orlando, Florida, December 2012. (AR: 44/231 = 19%)
92. Li Qian, Edmond W. W. Chan, Patrick P. C. Lee, and Cheng He. “Characterization of 3G Control-Plane Signaling Overhead from a Data-Plane Perspective.” Proceedings of the 15th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (**MSWiM 2012**), Paphos, Cyprus Island, October 2012. (AR: 42/161 = 26%)
93. Henry C. H. Chen¹ and Patrick P. C. Lee. “Enabling Data Integrity Protection in Regenerating-Coding-Based Cloud Storage.” Proceedings of the 31st IEEE International Symposium on Reliable Distributed Systems (**SRDS 2012**) (Full paper), Irvine, CA, October 2012. (AR: 28/114 = 24.6%)
94. Min Zheng, Patrick P. C. Lee, and John C. S. Lui. “ADAM: An Automatic and Extensible Platform to Stress Test Android Anti-Virus Systems”. Proceedings of the 9th Conference on Detection of Intrusions and Malware & Vulnerability Assessment (**DIMVA 2012**), Heraklion, Crete, Greece, July 2012. (AR: 14/44 = 31.8%)

95. Yunfeng Zhu³, Patrick P. C. Lee, Liping Xiang, Yinlong Xu and Lingling Gao. "A Cost-based Heterogeneous Recovery Scheme for Distributed Storage Systems with RAID-6 Codes". Proceedings of the 42nd IEEE/IFIP International Conference on Dependable Systems and Networks - Performance and Dependability Symposium (**DSN-PDS 2012**), Boston, MA, June 2012. (AR: 24/80 = 30%)
96. Xiuqiang He, Patrick P. C. Lee, Lujia Pan, Cheng He, and John C. S. Lui. "A Panoramic View of 3G Data/Control-Plane Traffic: Mobile Device Perspective." Proceedings of IFIP/TC6 Networking (**Networking 2012**), Prague, Czech Republic, May 2012. (AR: 64/225 = 28.4%)
97. Yunfeng Zhu³, Patrick P. C. Lee, Yuchong Hu, Liping Xiang, and Yinlong Xu. "On the Speedup of Single-Disk Failure Recovery in XOR-Coded Storage Systems: Theory and Practice." Proceedings of the 28th IEEE Conference on Mass Storage Systems and Technologies (**MSST 2012**), Monterey, CA, April 2012. (AR: 14/57 = 24.6%)
98. Chung-Pan Tang¹, Tsz-Yeung Wong, and Patrick P. C. Lee. "CloudVS: Enabling Version Control for Virtual Machines in an Open-Source Cloud under Commodity Settings." Proceedings of the 13th IEEE/IFIP Network Operations and Management Symposium (**NOMS 2012**), April 2012. (AR: 55/210 = 26.2%)
99. Yuchong Hu², Henry C. H. Chen¹, Patrick P. C. Lee, and Yang Tang¹. "NCCloud: Applying Network Coding for the Storage Repair in a Cloud-of-Clouds." Proceedings of the 10th USENIX Conference on File and Storage Technologies (**FAST 2012**) (Short paper), San Jose, CA, February 2012. (AR: 26/137 = 19.0%, short paper track = 7/38) (*The short paper track is held at the same quality bar as the full paper track.*)
100. Chun-Ho Ng¹, Mingcao Ma¹, Tsz-Yeung Wong, Patrick P. C. Lee and John C. S. Lui. "Live Deduplication Storage of Virtual Machine Images in an Open-Source Cloud." Proceedings of ACM/IFIP/USENIX 12th International Middleware Conference (**Middleware 2011**), Lisbon, Portugal, December 2011. (AR: 24/125 = 19.2%)
101. Ben S. Y. Fung¹ and Patrick P. C. Lee. "A Privacy-Preserving Defense Mechanism Against Request Forgery Attacks." Proceedings of the 10th IEEE International Conference on Trust, Security and Privacy in Computing and Communications (**TrustCom 2011**), November 2011. (AR: 126/471 = 26.8%) **Best paper award recipient.**
102. Harry F. W. Choi¹ and Patrick P. C. Lee. "An Extensible Design of a Load-Aware Virtual Router Monitor in User Space." 7th International Workshop on Scheduling and Resource Management for Parallel and Distributed Systems (**SRMPDS**) (in conjunction with ICPP 2011), Taipei, Taiwan, September 2011.
103. Arthur Rahumed¹, Henry C. H. Chen¹, Yang Tang¹, Patrick P. C. Lee, and John C. S. Lui. "A Secure Cloud Backup System with Assured Deletion and Version Control." 3rd International Workshop on Security in Cloud Computing (**CloudSec**) (in conjunction with ICPP 2011), Taipei, Taiwan, September 2011.
104. Qi Li³, Mingwei Xu, Jianping Wu, Patrick P. C. Lee, and Ke Xu. "An Algebraic Approach to Computing the Reliability of Internet Routing." Proceedings of IEEE International Conference on Computer Communications and Networks (**ICCCN 2011**), Maui, Hawaii, July 2011. (AR: 134/452 = 29.6%)
105. Yuchong Hu², Chiu-Man Yu², Yan Kit Li¹, Patrick P. C. Lee, and John C. S. Lui. "NCFS: On the Practicality and Extensibility of a Network-Coding-Based Distributed File System." Proceedings of the 2011 International Symposium on Network Coding (**NETCOD 2011**), Beijing, China, July 2011.
106. Patrick P. C. Lee, Hongying Liu, John C.S. Lui. "Analyzing Credit Evolution for Credit-based Incentive Schemes in Wireless Mesh Networks." Proceedings of the 8th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (**SECON 2011**), Salt Lake City, Utah, USA, June 2011. (AR: 66/241 = 27.4%)
107. Jin Cao, Aiyou Chen, and Patrick P. C. Lee. "Modeling Time Correlation in Passive Network Loss Tomography." Proceedings of the 41st IEEE/IFIP International Conference on Dependable Systems and Networks - Performance and Dependability Symposium (**DSN-PDS 2011**), Hong Kong, June 2011. (AR: 26/106 = 24.5%)

108. Mengwei Lu, Patrick P. C. Lee and John C. S. Lui. "Identity Attack and Anonymity Protection for P2P-VoD Systems." Proceedings of the ACM/IEEE International Workshop on Quality of Service (**IWQoS 2011**), San Jose, CA, USA, June 2011. (AR: 23/80 = 28.8%)
109. Qi Li³, Mingwei Xu, Jianping Wu, Xinwen Zhang, Patrick P. C. Lee and Ke Xu "Enhancing the Trust of Internet Routing with Lightweight Route Attestation." Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security (**ASIACCS 2011**), Hong Kong, March 2011. (AR: 35/217 = 16%)
110. Yang Tang¹, Patrick P. C. Lee, John C. S. Lui, and Radia Perlman "FADE: Secure Overlay Cloud Storage with File Assured Deletion." Proceedings of the 6th International ICST Conference on Security and Privacy in Communication Networks (**SecureComm 2010**), Singapore, September 2010. (AR: 28/112 = 25%)
111. Qi Li³, Mingwei Xu, Jianping Wu, Xingang Shi, Dah Ming Chiu, Patrick P.C. Lee. "Achieving Unified Protection for IP Routing." Proceedings of the 19th International Conference on Computer Communications and Networks (**ICCCN 2010**), Zurich, Switzerland, August 2010. (AR: 64/189 = 33.9%)
112. Patrick P. C. Lee, Tian Bu, Girish Chandranmenon. "A Lock-Free, Cache-Efficient Multi-Core Synchronization Mechanism for Line-Rate Network Traffic Monitoring." Proceedings of 24th IEEE International Parallel and Distributed Processing Symposium (**IPDPS 2010**), Atlanta, Georgia, April 2010. (AR: 127/527 = 24%)
113. Patrick P. C. Lee, Vishal Misra, and Dan Rubenstein. "On the Robustness of Wireless Opportunistic Routing Toward Inaccurate Link-Level Measurements." Proceedings of the 2nd International Conference on Communication Systems and Networks (**COMSNETS**) (invited paper), Bangalore, India, January 2010.
114. Patrick P. C. Lee, Tian Bu, and Girish Chandranmenon. "A Lock-Free, Cache-Efficient Shared Ring Buffer for Multi-Core Architectures." Proceedings of ACM/IEEE Symposium on Architectures for Networking and Communications Systems (**ANCS'09**) (Poster session), Princeton, New Jersey, October 2009.
115. Paramvir (Victor) Bahl, Ranveer Chandra, Patrick P. C. Lee, Vishal Misra, Jitendra Padhye, Dan Rubenstein, and Yan Yu. "Opportunistic Use of Client Repeaters to Improve Performance of WLANs." Proceedings of the 4th ACM International Conference on emerging Networking Experiments and Technologies (**CoNEXT'08**), Madrid, Spain, December 2008. (AR: 29/166 = 17%) **Best paper award recipient.**
116. Tian Bu, Jin Cao, Aiyu Chen, and Patrick P. C. Lee. "A Fast and Compact Method for Unveiling Significant Patterns in High Speed Networks." Proceedings of the 26th Annual IEEE Conference on Computer Communications (**INFOCOM'07**), Anchorage, Alaska, May 2007. (AR: 252/1400 = 18%)
117. Patrick P. C. Lee, Tian Bu, and Thomas Woo. "On the Detection of Signaling DoS Attacks on 3G Wireless Networks." Proceedings of the 26th Annual IEEE Conference on Computer Communications (**INFOCOM'07**), Anchorage, Alaska, May 2007. (AR: 252/1400 = 18%)
118. Patrick P. C. Lee, Vishal Misra, and Dan Rubenstein. "Toward Optimal Network Fault Correction via End-to-End Inference." Proceedings of the 26th Annual IEEE Conference on Computer Communications (**INFOCOM'07**), Anchorage, Alaska, May 2007. (AR: 252/1400 = 18%)
119. Patrick P. C. Lee, Vishal Misra, and Dan Rubenstein. "Distributed Algorithms for Secure Multipath Routing." Proceedings of the 24th Annual IEEE Conference on Computer Communications (**INFOCOM'05**), Miami, FL, March 2005. (AR: 244/1419 = 17%)
120. Patrick P. C. Lee, John C. S. Lui, and David K. Y. Yau. "Distributed Collaborative Key Agreement Protocols for Dynamic Peer Groups." In Proceedings of the 10th IEEE International Conference on Network Protocols (**ICNP'02**), November 2002. (AR: 32/217 = 14.7%)

2.1.2 Journal Publications

1. Haifeng Sun, Qun Huang, Patrick P. C. Lee, Wei Bai, Feng Zhu, and Yungang Bao. "Distributed Network Telemetry with Resource Efficiency and Full Accuracy." Accepted for publication in **IEEE/ACM Transactions on Networking (TON)**. (An earlier version appeared in SIGCOMM 2020)

2. Si Wu, Zhirong Shen, Patrick P. C. Lee, Zhiwei Bai, and Yinlong Xu. “Elastic Reed-Solomon Codes for Efficient Redundancy Transitioning in Distributed Key-Value Stores.” Accepted for publication in **IEEE/ACM Transactions on Networking (TON)**.
3. Lu Tang, Qun Huang, and Patrick P. C. Lee. “MVPipe: Enabling Lightweight Updates and Fast Convergence in Hierarchical Heavy Hitter Detection.” Accepted for publication in **IEEE/ACM Transactions on Networking (TON)**.
4. Qiang Zhang, Yongkun Li, Patrick P. C. Lee, and Yinlong Xu. “The Design and Implementation of UniKV for Mixed Key-Value Storage Workloads.” Accepted for publication in **IEEE Transactions on Knowledge and Data Engineering (TKDE)**.
5. Guowen Gong, Zhirong Shen, Liang Chen, Suzhen Wu, Xiaolu Li, Patrick P. C. Lee, Zhiguo Wan, and Jiwu Shu. “Optimal Rack-Coordinated Updates in Erasure-Coded Data Centers: Design and Analysis.” **IEEE Transactions on Computers (TC)**, 72(7), pp. 1871-1885, July 2023.
6. Jinhong Li¹, Qiuping Wang¹, Patrick P. C. Lee, and Chao Shi. “An In-Depth Comparative Analysis of Cloud Block Storage Workloads: Findings and Implications.” **ACM Transactions on Storage (TOS)**, 19(2), pp. 16:1-16:32, May 2023.
7. Lu Tang, Yao Xiao, Qun Huang, and Patrick P. C. Lee. “A High-Performance Invertible Sketch for Network-Wide Superspreader Detection.” **IEEE/ACM Transactions on Networking (TON)**, 31(2), pp.724-737, April 2023.
8. Siyuan Sheng¹, Qun Huang, and Patrick P. C. Lee. “A General Delta-based In-band Network Telemetry Framework with Extremely Low Bandwidth Overhead.” **Elsevier Computer Networks (COMNET)**, 223, March 2023.
9. Shujie Han¹, Patrick P. C. Lee, Zhirong Shen, Cheng He, Yi Liu, and Tao Huang. “StreamDFP: A General Stream Mining Framework for Adaptive Disk Failure Prediction.” **IEEE Transactions on Computers (TC)**, 72(2), pp. 520-534, February 2023.
10. Hanxu Hou, Yunghsiang S. Han, Patrick P. C. Lee, You Wu, Guojun Han, Mario Blaum. “A Generalization of Array Codes With Local Properties and Efficient Encoding/Decoding.” **IEEE Transactions on Information Theory (TIT)**, 69(1), pp. 107-125, January 2023.
11. Xiaolu Li, Keyun Cheng¹, Zhirong Shen, and Patrick P. C. Lee. “Fast Proactive Repair in Erasure-Coded Storage: Analysis, Design, and Implementation.” **IEEE Transactions on Parallel and Distributed Systems (TPDS)**, 33(12), pp. 3400-3414, December 2022.
12. Zuoru Yang¹, Jingwei Li, Yanjing Ren, and Patrick P. C. Lee. “Tunable Encrypted Deduplication with Attack-Resilient Key Management.” **ACM Transactions on Storage (TOS)**, 18(4), pp. 32:1-32:38, November 2022.
13. Mi Zhang¹, Qiuping Wang¹, Zhirong Shen, and Patrick P. C. Lee. “POCache: Toward Robust and Congurable Straggler Tolerance with Parity-Only Caching.” **Elsevier Journal of Parallel and Distributed Computing (JPDC)**, 167, pp. 157-172, September 2022.
14. Jingwei Li, Suyu Huang, Yanjing Ren, Zuoru Yang¹, Patrick P. C. Lee, Xiaosong Zhang, and Yao Hao. “Enabling Secure and Space-Efficient Metadata Management in Encrypted Deduplication.” **IEEE Transactions on Computers (TC)**, 71(4), pp. 959-970, April 2022.
15. Yuchong Hu, Xiaoyang Zhang, Patrick P. C. Lee, and Pan Zhou. “NCScale: Toward Optimal Storage Scaling via Network Coding.” **IEEE/ACM Transactions on Networking (TON)**, 30(1), pp. 271-284, February 2022.
16. Si Wu, Zhirong Shen, Patrick P. C. Lee, Yinlong Xu. “Optimal Repair-Scaling Trade-off in Locally Repairable Codes: Analysis and Evaluation.” **IEEE Transactions on Parallel and Distributed Systems (TPDS)**, 33(1), pp. 56-69, January 2022.

17. Xiaolu Li, Zuoru Yang, Jinhong Li, Runhui Li, Patrick P. C. Lee, Qun Huang, and Yuchong Hu. "Repair Pipelining for Erasure-Coded Storage: Algorithms and Evaluation." **ACM Transactions on Storage (TOS)**, 17(2), 13:1-13:29, May 2021.
18. Hanxu Hou, Yunghsiang S. Han, and Patrick P. C. Lee. "Two Classes of Binary MDS Array Codes with Asymptotically Optimal Repair for Any Single Column." **IEEE Transactions on Communications (TCOM)**, 68(11), pp. 6723-6736, November 2020.
19. Lu Tang¹, Qun Huang, and Patrick P. C. Lee. "A Fast and Compact Invertible Sketch for Network-Wide Heavy Flow Detection." **IEEE/ACM Transactions on Networking (TON)**, 28(5), pp. 2350-2363, October 2020.
20. Hanxu Hou, Patrick P. C. Lee, and Yunghsiang S. Han. "Zigzag-Decodable Reconstruction Codes with Asymptotically Optimal Repair for All Nodes." **IEEE Transactions on Communications (TCOM)**, 68(10), pp. 5999-6011, October 2020.
21. Zhirong Shen² and Patrick P. C. Lee. "Cross-Rack-Aware Updates in Erasure-Coded Data Centers: Design and Evaluation." **IEEE Transactions on Parallel and Distributed Systems (TPDS)**, 31(10), pp. 2315-2328, October 2020.
22. Zhirong Shen², 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)**, 17(2), pp. 248-261, March/April 2020.
23. Jingwei Li, Patrick P. C. Lee, Chufeng Tan, Chuan Qin¹, and Xiaosong Zhang. "Information Leakage in Encrypted Deduplication via Frequency Analysis: Attacks and Defenses." **ACM Transactions on Storage (TOS)**, 16(1), pp. 4:1-4:30, March 2020.
24. Hanxu Hou and Patrick P. C. Lee. "Binary MDS Array Codes with Optimal Repair." **IEEE Transactions on Information Theory (TIT)**, 66(3), pp. 1405-1422, March 2020.
25. Lujia Pan, Jianfeng Zhang, Patrick P. C. Lee, Marcus Kalander, Junjian Ye, and Pinghui Wang. "Proactive Microwave Link Anomaly Detection in Cellular Data Networks." **Elsevier Computer Networks (COMNET)**, 167, February 2020.
26. Mi Zhang¹, Shujie Han¹, and Patrick P. C. Lee. "SimEDC: A Simulator for the Reliability Analysis of Erasure-Coded Data Centers." **IEEE Transactions on Parallel and Distributed Systems (TPDS)**, 30(12), pp. 2836-2848, December 2019.
27. Qi Li, Xiaoyue Zou, Qun Huang, Jing Zheng, and Patrick P.C. Lee. "Dynamic Packet Forwarding Verification in SDN." **IEEE Transactions on Dependable and Secure Computing (TDSC)**, 16(6), pp. 915-929, November/December 2019.
28. Hanxu Hou, Yunghsiang S. Han, Patrick P. C. Lee, Yuchong Hu, and Hui Li. "A New Design of Binary MDS Array Codes with Asymptotically Weak-Optimal Repair." **IEEE Transactions on Information Theory (TIT)**, 65(11), pp. 7095-7113, November 2019.
29. Zhinan Cheng¹, Qun Huang, and Patrick P. C. Lee. "On the Performance and Convergence of Distributed Stream Processing via Approximate Fault Tolerance." **The VLDB Journal (VLDBJ)**, 28(5), pp. 821-846, October 2019.
30. Yongkun Li, Helen H. W. Chan¹, Patrick P. C. Lee, and Yinlong Xu. "Enabling Efficient Updates in KV Storage via Hashing: Design and Performance Evaluation." **ACM Transactions on Storage (TOS)**, 15(3), pp. 20:1-20:29, August 2019.
31. Hanxu Hou, Patrick P. C. Lee, Kenneth W. Shum, and Yuchong Hu. "Rack-Aware Regenerating Codes for Data Centers." **IEEE Transactions on Information Theory (TIT)**, 65(8), pp. 4730-4745, August 2019.
32. Zhirong Shen², 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), pp. 1552-1564, July 2019.

33. Zhirong Shen², 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), pp. 2757-2770, December 2018.
34. Helen H. W. Chan¹, Yongkun Li, Patrick P. C. Lee, and Yinlong Xu. “Elastic Parity Logging for SSD RAID Arrays: Design, Analysis, and Implementation.” **IEEE Transactions on Parallel and Distributed Systems (TPDS)**, 29(10), pp. 2241-2253, October 2018.
35. Min Fu³, Shujie Han¹, Patrick P. C. Lee, Dan Feng, Zuoning Chen, and Yu Xiao. “A Simulation Analysis of Redundancy and Reliability in Primary Storage Deduplication.” **IEEE Transactions on Computers (TC)**, 67(9), pp. 1259-1272, September 2018.
36. Qi Li, Yanyu Chen, Patrick P. C. Lee, Kui Ren, and Mingwei Xu. “Security Policy Violations in SDN Data Plane.” **IEEE/ACM Transactions on Networking (TON)**, 26(4), pp. 1715-1727, August 2018.
37. Mingqiang Li², Runhui Li¹, and Patrick P. C. Lee. “Relieving Both Storage and Recovery Burdens in Big Data Clusters with R-STAIR Codes.” **IEEE Internet Computing**, 22(4), pp. 15-26, July-August 2018.
38. Hanxu Hou and Patrick P. C. Lee. “A New Construction of EVENODD Codes with Lower Computational Complexity.” **IEEE Communications Letters (CL)**, 22(6), pp. 1120-1123, June 2018.
39. Yuchong Hu, Patrick P. C. Lee, Kenneth W. Shum, and Pan Zhou. “Proxy-assisted Regenerating Codes with Uncoded Repair for Distributed Storage Systems.” **IEEE Transactions on Information Theory (TIT)**, 64(4), pp. 2512-2528, April 2018.
40. Yuchong Hu, Xiaolu Li¹, Mi Zhang¹, Patrick P. C. Lee, Xiaoyang Zhang, Pan Zhou, and Dan Feng. “Optimal Repair Layering for Erasure-Coded Data Centers: From Theory to Practice.” **ACM Transactions on Storage (TOS)**, 13(4), pp. 33:1-33:24, November 2017.
41. Peng Zhang, Qi Li, and Patrick P. C. Lee. “Achieving Content-Oriented Anonymity with CRISP.” **IEEE Transactions on Dependable and Secure Computing (TDSC)**, 14(6), pp. 578-590, November-December, 2017.
42. Qi Li, Patrick P. C. Lee, Peng Zhang, Purui Su, Liang He, and Kui Ren. “Capability-Based Security Enforcement in Named Data Networking.” **IEEE/ACM Transactions on Networking (TON)**, 25(5), pp. 2719-2730, October 2017.
43. Runhui Li¹, Yuchong Hu, and Patrick P. C. Lee. “Enabling Efficient and Reliable Transition from Replication to Erasure Coding for Clustered File Systems.” **IEEE Transactions on Parallel and Distributed Systems (TPDS)**, 28(9), pp. 2500-2513, September 2017.
44. Chuan Qin¹, Jingwei Li, and Patrick P. C. Lee. “The Design and Implementation of a Rekeying-aware Encrypted Deduplication Storage System.” **ACM Transactions on Storage (TOS)**, 13(1), 9:1-9:30, March 2017.
45. 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), pp. 877-890, March 2017.
46. Yaliang Li, Jing Gao, Patrick P. C. Lee, Lu Su, Caifeng He, Cheng He, Fan Yang, and Wei Fan. “A Weighted Crowdsourcing Approach for Network Quality Measurement in Cellular Data Networks.” **IEEE Transaction on Mobile Computing (TMC)**, 16(2), pp. 300-313, February 2017.
47. Qun Huang¹ and Patrick P. C. Lee. “Toward High-Performance Distributed Stream Processing via Approximate Fault Tolerance.” Proceedings of the VLDB Endowment (**PVLDB**), 10(3), pp. 73-84, November 2016 (Also presented at VLDB 2017, Munich, Germany).
48. Mingqiang Li², Chuan Qin¹, Jingwei Li², and Patrick P. C. Lee. “CDStore: Toward Reliable, Secure, and Cost-Efficient Cloud Storage via Convergent Dispersal.” **IEEE Internet Computing**, 20(3), pp. 45-53, May-June 2016 (Special issue: Cloud Storage).

49. Yongkun Li², Patrick P. C. Lee, and John C. S. Lui. "Analysis of Reliability Dynamics of SSD RAID." **IEEE Transaction on Computers (TC)**, 65(4), pp. 1131-1144, April 2016 (Special issue: Emerging Memory Technologies).
50. Qun Huang¹ and Patrick P. C. Lee. "A Hybrid Local and Distributed Sketching Design for Accurate and Scalable Heavy Key Detection in Network Data Streams." **Computer Networks (COMNET)**, Elsevier, 91, pp. 298-315, November 2015.
51. Yunfeng Zhu³, Jian Lin¹, Patrick P. C. Lee, and Yinlong Xu. "Boosting Degraded Reads in Heterogeneous Erasure-Coded Storage Systems." **IEEE Transactions on Computers (TC)**, 64(8), pp. 2145-2157, August 2015.
52. Runhui Li¹, Jian Lin¹, and Patrick P. C. Lee. "Enabling Concurrent Failure Recovery for Regenerating-Coding-Based Storage Systems: From Theory to Practice." **IEEE Transactions on Computers (TC)**, 64(7), pp. 1898-1911, July 2015.
53. Jin Li, Yan-Kit Li¹, Xiaofeng Chen, Patrick P. C. Lee, and Wenjing Lou. "A Hybrid Cloud Approach for Secure Authorized Deduplication." **IEEE Transactions on Parallel and Distributed Systems (TPDS)**, 26(5), pp. 1206-1216, May 2015.
54. Yan-Kit Li¹, Min Xu¹, Chun-Ho Ng¹, and Patrick P. C. Lee. "Efficient Hybrid Inline and Out-of-line Deduplication for Backup Storage." **ACM Transactions on Storage (TOS)**, 11(1), pp. 2:1-2:21, February 2015.
55. Chung Pan Tang¹, Patrick P. C. Lee, and Tsz-Yeung Wong. "Tunable Version Control System for Virtual Machines in an Open-Source Cloud." **IEEE Transactions on Services Computing (TSC)**, 8(1), pp. 155-168, January-February 2015.
56. Mingqiang Li² and Patrick P. C. Lee. "STAIR Codes: A General Family of Erasure Codes for Tolerating Device and Sector Failures." **ACM Transactions on Storage (TOS)**, 10(4), pp. 14:1-14:30, October 2014. (One of FAST 2014 papers selected for fast-track publication in TOS)
57. Zhuo Zhang, Zhibin Zhang, Patrick P. C. Lee, Yunjie Liu, and Gaogang Xie. "Toward Unsupervised Protocol Feature Word Extraction." **IEEE Journal on Selected Areas in Communications (JSAC)**, 32(10), pp. 1894-1906, October 2014.
58. Yunfeng Zhu³, Patrick P. C. Lee, Yinlong Xu, Yuchong Hu, and Liping Xiang "On the Speedup of Recovery in Large-Scale Erasure-Coded Storage Systems." **IEEE Transactions on Parallel and Distributed Systems (TPDS)**, 25(7), pp. 1830-1840, July 2014.
59. Jin Li, Xiaofeng Chen, Mingqiang Li², Jingwei Li, Patrick P. C. Lee, and Wenjing Lou. "Secure Deduplication with Efficient and Reliable Convergent Key Management." **IEEE Transactions on Parallel and Distributed Systems (TPDS)**, 25(6), pp. 1615-1625, June 2014.
60. Yongkun Li², Patrick P. C. Lee, John C. S. Lui. "Stochastic Modeling and Optimization of Garbage Collection Algorithms in Solid State Drive Systems." **Queueing Systems: Theory and Applications (QUESTA)**, Springer, 77(2), pp. 115-148, June 2014.
61. Silei Xu, Runhui Li¹, Patrick P. C. Lee, Yunfeng Zhu³, Liping Xiang, Yinlong Xu, and John C. S. Lui. "Single Disk Failure Recovery for X-code-based Parallel Storage Systems." **IEEE Transactions on Computers (TC)**, 63(4), pp. 995-1007, April 2014.
62. Henry C. H. Chen¹ and Patrick P. C. Lee. "Enabling Data Integrity Protection in Regenerating-Coding-Based Cloud Storage: Theory and Implementation." **IEEE Transactions on Parallel and Distributed Systems (TPDS)**, 25(2), pp. 407-416, February 2014 (Special Issue: Trust, Security, and Privacy, AR: 24/150 = 16%).
63. Henry C. H. Chen¹, Yuchong Hu², Patrick P. C. Lee, and Yang Tang¹. "NCCloud: A Network-Coding-Based Storage System in a Cloud-of-Clouds." **IEEE Transactions on Computers (TC)**, 63(1), pp. 31-44, January 2014 (Special Issue: Cloud of Clouds).

64. Hongying Liu, Patrick P. C. Lee, John C. S. Lui. "On the Credit Evolution of Credit-Based Incentive Protocols in Wireless Mesh Networks." **Computer Networks (COMNET)**, Elsevier, 57(17), pp. 3327-3343, December 2013.
65. Qun Huang¹ and Patrick P. C. Lee. "An Experimental Study of Cascading Performance Interference in a Virtualized Environment." **ACM SIGMETRICS Performance Evaluation Review (PER)** (invited paper), 40(4), pp. 43-52, March 2013.
66. Yang Tang¹, Patrick P. C. Lee, John C. S. Lui, and Radia Perlman "Secure Overlay Cloud Storage with Access Control and Assured Deletion." **IEEE Transactions on Dependable and Secure Computing (TDSC)**, 9(6), pp. 903-916, November-December 2012.
67. Peng Zhang³, Yixin Jiang, Chuang Lin, Patrick P. C. Lee, and John C. S. Lui "ANOC: Anonymous Network-Coding-Based Communication with Efficient Cooperation." **IEEE Journal on Selected Areas in Communications (JSAC)**, 30(9), pp. 1738-1745, October 2012.
68. Mingwei Xu, Qi Li³, Patrick P. C. Lee, Yanhai Peng, and Jianping Wu. "VegaNet: A Virtualized Experimentation Platform for Production Networks with Connectivity Consistency." **IEEE Network Magazine**, 26(5), pp. 15-21, September-October 2012 (Special Issue on Wired and Wireless Network Virtualization).
69. Qi Li³, Mingwei Xu, Jianping Wu, Patrick P. C. Lee, Xingang Shi, Dah Ming Chiu, and Yuan Yang. "A Unified Approach to Routing Protection in IP Networks." **IEEE Transactions on Network and Service Management (TNSM)**, 9(3), pp. 306-319, September 2012.
70. Qi Li³, Mingwei Xu, Jianping Wu, Xinwen Zhang, Patrick P. C. Lee, and Ke Xu "Enhancing the Trust of Internet Routing with Lightweight Route Attestation." **IEEE Transactions on Information Forensics and Security (TIFS)**, 7(2), pp. 691-703, April 2012.
71. Qi Li³, Mingwei Xu, Jianping Wu, Patrick P. C. Lee, and Dah Ming Chiu "Towards a Practical Approach for BGP Stability with Root Cause Check." **Journal of Parallel and Distributed Computing (JPDC)**, Elsevier, 71(8), pp. 1098-1110, August 2011.
72. Yan Cai, Patrick P. C. Lee, Weibo Gong, and Don Towsley "Analysis of Traffic Correlation Attacks on Router Queues." **Computer Networks (COMNET)**, Elsevier, 55(3), pp. 734-747, February 2011.
73. Tian Bu, Jin Cao, Aiyu Chen, and Patrick P. C. Lee. "Sequential Hashing: A Flexible Approach for Unveiling Significant Patterns in High Speed Networks." **Computer Networks (COMNET)**, Elsevier, 54(18), pp. 3309-3326, December 2010.
74. Patrick P. C. Lee. "Assured Deletion of Digital Files on Cloud." **IEEE COMSOC MMTC E-Letter (invited paper)**, 5(6), pp. 43-44, November 2010.
75. Patrick P. C. Lee, Vishal Misra, and Dan Rubenstein. "Toward Optimal Network Fault Correction in Externally Managed Overlay Networks." **IEEE Transactions on Parallel and Distributed Systems (TPDS)**, 21(3), pp. 354-366, March 2010.
76. Patrick P. C. Lee, Tian Bu, and Thomas Woo. "On the Detection of Signaling DoS Attacks on 3G/WiMax Wireless Networks." **Computer Networks (COMNET)**, 53(15), pp. 2601-2616, October 2009.
77. Paramvir (Victor) Bahl, Ranveer Chandra, Patrick P. C. Lee, Vishal Misra, Jitendra Padhye, Dan Rubenstein, and Yan Yu. "Opportunistic Use of Client Repeaters to Improve Performance of WLANs." **IEEE/ACM Transactions on Networking (ToN)**, 17(4), pp. 1160-1171, August 2009.
78. Patrick P. C. Lee, Vishal Misra, and Dan Rubenstein. "Distributed Algorithms for Secure Multipath Routing in Attack-Resistant Networks." **IEEE/ACM Transactions on Networking (ToN)**, 15(6), pp. 1490-1501, December 2007.
79. Patrick P. C. Lee, John C. S. Lui, and David K. Y. Yau. "SEAL: A Secure Communication Library for Building Dynamic Group Key Agreement Applications." **Journal of Systems and Software (JSS)**, 80(3), pp. 356-370, March 2007.

80. Patrick P. C. Lee, John C. S. Lui, and David K. Y. Yau. "Distributed Collaborative Key Agreement and Authentication Protocols for Dynamic Peer Groups." **IEEE/ACM Transactions on Networking (ToN)**, 14(2), pp. 263-276, April 2006.

2.1.3 Books

1. Cheng He, Mengling Feng, Patrick P. C. Lee, Pinghui Wang, Shujie Han¹, Yi Liu (Eds.). "Large-Scale Disk Failure Prediction", pp. 1-143, Springer, June 2020 (ISBN: 978-981-15-7749-9).

2.1.4 Preprints/Technical Reports

1. Shujie Han¹, Jun Wu, Erci Xu, Cheng He, Patrick P. C. Lee, Yi Qiang, Qixing Zheng, Tao Huang, Zixi Huang, and Rui Li. "Robust Data Preprocessing for Machine-Learning-Based Disk Failure Prediction in Cloud Production Environments." arXiv:1912.09722, December 2019.
2. Rong Gu, Qianhao Dong, Haoyuan Li, Joseph Gonzalez, Zhao Zhang, Shuai Wang, Yihua Huang, Scott Shenker, Ion Stoica, and Patrick P. C. Lee. "DFS-Perf: A Scalable and Unified Benchmarking Framework for Distributed File Systems." EECS Department, University of California, Berkeley, Technical Report No. UCB/EECS-2016-133, July 2016.

2.2 Research Awards/Recognitions

1. Young Researcher Award, The Chinese University of Hong Kong, 2020-21
2. Prof. C. V. Ramamoorthy Best Paper Award, "Enabling I/O-Efficient Redundancy Transitioning in Erasure-Coded KV Stores via Elastic Reed-Solomon Codes", 39th International Symposium on Reliable Distributed Systems (SRDS), 2020 (1 out of 131 submissions)
3. Best paper award, "Multi-Layer Transformed MDS Codes with Optimal Repair Access and Low Sub-Packetization", The Chinese Institute of Electronics, China Academic Annual Conference on Information Theory (CIEIT), 2019
4. Best paper award, "Triple-Fault-Tolerant Binary MDS Array Codes with Asymptotically Optimal Repair", The Chinese Institute of Electronics, China Academic Annual Conference on Information Theory (CIEIT), 2017
5. Distinguished Collaborator Award, Huawei Noah's Ark Lab, February 2017
6. Best paper award nominee, "Information Leakage in Encrypted Deduplication via Frequency Analysis", IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), 2017 (3 out of 220 submissions)
7. Best paper award, "A Privacy-Preserving Defense Mechanism Against Request Forgery Attacks", IEEE TrustCom 2011 (2 out of 471 submissions)
8. Best paper award, "Opportunistic Use of Client Repeaters to Improve Performance of WLANs", ACM CoNEXT 2008 (1 out of 166 submissions)

2.3 Professional Activities

2.3.1 Journal Editorship

1. Associate Editor, ACM Transactions on Storage, August 2018 - present.
2. Associate Editor, IEEE/ACM Transactions on Networking, October 2017 - present.
3. Guest Editor, IEEE Transactions on Cloud Computing: Special Issue on Cloud Security, 2(4), Oct-Dec 2014.
4. Editor, Elsevier Computer Standards and Interfaces, 2010-2014

2.3.2 Conference TPC

TPC Chairmanship:

1. PC co-chair, 55th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (**DSN 2025**), June 2025.
2. PC co-chair, 28th IEEE International Symposium on Local and Metropolitan Area Networks (**LANMAN 2022**), July 2022.
3. PC co-chair, ACM Asia-Pacific Workshop on Systems (**APSys 2020**), August 2020.

TPC Membership:

1. ACM European Conference on Computer Systems (**EuroSys**), 2025, 2022 (light PC), 2021 (light PC)
2. IEEE International Conference on Computer Communications (**INFOCOM**), 2024, 2023, 2017, 2016, 2015, 2014, 2013
3. ACM/SPEC International Conference on Performance Engineering (**ICPE**), 2024, 2022, 2019, 2018, 2017
4. 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (**ASPLOS**), 2024
5. IEEE/IFIP International Conference on Dependable Systems and Networks (**DSN**), 2024 (Disrupt Track), 2020, 2017, 2016, 2013
6. European Dependable Computing Conference (**EDCC**), 2024, 2022, 2021, 2020, 2019, 2018
7. USENIX Conference on File and Storage Technologies (**FAST**), 2024, 2023, 2022, 2021, 2020
8. IEEE International Conference on Network Protocols (**ICNP**), 2024, 2023, 2022, 2021, 2020, 2019, 2018, 2017, 2016, 2015, 2014, 2013, 2011, 2010
9. IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (**CCGrid**), 2024, 2022, 2021, 2019, 2018
10. ACM International Conference on Emerging Networking Experiments and Technologies (**CoNEXT**), 2023, 2013, 2008 (Shadow PC)
11. International Conference on Algorithms and Architectures for Parallel Processing (ICA3PP), 2023 (track co-chair)
12. IEEE International Conference on Distributed Computing Systems (**ICDCS**), 2023 (track co-chair), 2022, 2021, 2018, 2017, 2012
13. ACM International Storage and Systems Conference (**SYSTOR**), 2023, 2021, 2020, 2019
14. IEEE/ACM International Symposium on Quality of Service (**IWQoS**), 2023, 2021, 2020, 2014, 2012, 2011, 2010
15. IEEE Non-Volatile Memory Systems and Applications Symposium (**NVMSA**), 2022, 2021, 2020, 2019, 2018, 2017, 2016, 2015, 2014
16. Workshop on Hot Topics in Storage and File Systems (**HotStorage**), 2022
17. IEEE International Symposium on Local and Metropolitan Area Networks (**LANMAN**), 2021, 2020, 2019, 2018
18. IEEE International Symposium on Reliable Distributed Systems (**SRDS**), 2020, 2017
19. IFIP Networking (**Networking**), 2020, 2019, 2010
20. USENIX Symposium on Networked Systems Design and Implementation (**NSDI**), 2020 (light PC)
21. USENIX Annual Technical Conference (**ATC**), 2019
22. IEEE International Symposium on High Assurance Systems Engineering (**HASE**), 2019, 2016, 2015, 2014

23. IEEE International Conference on Algorithms and Architectures for Parallel Processing (**ICA3PP**), 2018
24. IEEE International Conference on Parallel and Distributed Systems (**ICPADS**), 2017, 2016
25. IEEE Conference on Dependable and Secure Computing (**DSC**), 2017
26. International Workshop of Software-Defined Data Communications and Storage (**SDDCS**), 2017, 2016, 2015
27. International Teletraffic Congress (**ITC**), 2016 (Area co-chair)
28. International Conference on Parallel Processing (**ICPP**), 2016
29. IEEE International Symposium on Multimedia (**ISM**), 2015
30. CSI International Symposium on Computer Architecture and Digital Systems (**CADS**), 2015
31. IEEE International Conference on Communications (**ICC**), 2015
32. International Conference on Communication System and Networks (**COMSNETS**), 2015
33. Information Security Conference (**ISC**), 2014
34. International Conference on Network and System Security (**NSS**), 2014
35. IEEE International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (**MASCOTS**), 2014
36. International Conference on Security and Privacy in Communication Networks (**SecureComm**), 2014, 2013
37. International Conference on Security and Cryptography (**SECURITY**), 2014, 2013, 2012
38. International Symposium on Cyberspace Safety and Security (**CSS**), 2014
39. IEEE International Conference on Networking, Architecture, and Storage (**NAS**), 2014
40. IEEE International Performance Computing and Communications Conference (**IPCCC**), 2013
41. Workshop on Distributed Storage Systems and Coding for BigData (in conjunction with IEEE BigData), 2013
42. ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization (**WiNTECH**), 2013, 2012, 2011, 2010, 2009
43. IEEE International Symposium on Network Coding (**NetCod**), 2013
44. IEEE Wireless Communications and Networking Conference (**WCNC**), 2013, 2009
45. International Conference on Computing, Networking and Communications (**ICNC**), 2013, 2012
46. IEEE International Symposium on Trust and Security in Cloud Computing (**TSCloud**), 2012, 2011
47. IEEE Vehicular Technology Conference (**VTC**), 2011 Fall, 2010 Fall, 2010 Spring
48. ACM Special Interest Group on Measurement and Evaluation (**SIGMETRICS**), 2011
49. IEEE International Conference on Computer Communication Networks (**ICCCN**), 2010
50. IEEE International Conference on Advanced Information Networking and Applications (**AINA**), 2010
51. ACM International Conference on Mobile Computing and Networking (**MOBICOM**), 2009 (Shadow PC)

2.3.3 Conference Organization

1. General co-chair, 29th IEEE International Symposium on Local and Metropolitan Area Networks (LANMAN 2023), July 2023.
2. Publication co-chair, 30rd IEEE International Conference on Network Protocols (ICNP 2022), November 2022
3. Publication co-chair, 29th IEEE International Conference on Network Protocols (ICNP 2021), November 2021
4. Registration co-chair, ACM Symposium on Operating Systems Principles (SOSP), 2017

5. Local arrangement co-chair, ACM SIGCOMM, August 2013
6. Publications chair, ACM SIGMETRICS, June 2010

2.4 US Patents Granted

1. US patent number: 11640257. “Data Processing Method and Apparatus.” Inventors: Cheng He, Qun Huang, Pak-Ching Lee. Filed in January 2019. Granted in 2023.
2. US patent number: 11228513. “Traffic Measurement Method, Device, and System.” Inventors: Qun Huang, Runhui Li, Yong Huang, Gong Zhang, Pak-Ching Lee, Lu Tang. Filed in January 2020. Granted in January 2022.
3. US patent number: 11003533. “Data Processing Method, System, and Apparatus.” Inventors: Jiajin Zhang, Matt M. T. Yiu, Pak-Ching Lee. Filed in March 2019. Granted in May 2021.
4. US patent number: 9923794. “Method, Apparatus, and System for Identifying Abnormal IP Data Stream.” Inventors: Cheng He, Qun Huang, Pak-Ching Lee. Filed in July 2015. Granted in March 2018.
5. US patent number: 8751737. “Method and apparatus for using a shared ring buffer to provide thread synchronization in a multi-core processor system.” Inventors: Tian Bu, Girish Chandranmenon, Pak-Ching Lee. Filed in June 2009. Granted in June 2014.
6. US patent number: 7779143. “Scalable methods for detecting significant traffic patterns in a data network.” Inventors: Tian Bu, Jin Cao, Aiyu Chen, Pak-Ching Lee. Filed in June 2007. Granted in August 2010.

2.5 Industrial Impact Beyond Publications

Some of my research work involves close collaboration with industry.

- **CU Coding.** I co-founded a company, CU Coding, in 2018 on commercializing dependable multi-cloud storage. We design **nEdge/nCloud**, a software-defined multi-cloud storage system that provides fault tolerance, security, storage efficiency, scalability, performance, and configurability guarantees. nEdge/nCloud builds on our prior research on network coding storage and secure data protection to interconnect the storage resources of multiple clouds. Our nEdge/nCloud project, named under Nexoedge, is the first standardized open-source university project from Hong Kong at Linux Foundation Edge (LF Edge).
- **Alibaba.** Since 2018, I have been collaborating with Alibaba to work on AI-based data center management and have received three funded projects. We design machine learning techniques for disk failure prediction in data centers. We also study the I/O workloads in the cloud block storage system deployed at Alibaba.
- **Huawei.** Since 2010, I collaborated with Huawei on applying AI and data management algorithms to telecommunication network management as well as designing practical regenerating codes for rack-based data centers. The collaboration leads to several publications and patents. I also received the Distinguished Collaborator Award from Huawei in February 2017.
- **Alcatel-Lucent.** In 2007-11, I collaborated with Alcatel-Lucent as an intern and a consultant on cellular network management. I was responsible for designing and implementing a multi-core traffic analyzer for 3G wireless networks. The analyzer supports protection against various network anomalies, multi-threaded packet analysis at line speed, and auditing of mobiles. It is part of the award-winning product Alcatel-Lucent 9900 Wireless Network Guardian.

3 Contributions in Teaching

3.1 Curriculum Development

I served as the chair of the Curriculum Committee since Fall 2017. I was responsible for overseeing the curriculum design to ensure that the offered courses provide students with sufficient trainings on computer science as well as reflect the latest trends of technologies. Two of my major tasks include:

- I took a leading role in developing the new undergraduate major programme *Artificial Intelligence: Systems and Technologies (AIST)*, which was launched in Fall 2019. The programme features a strong inter-disciplinary engineering curriculum that includes cohesive efforts across all departments in the Faculty of Engineering.
- I took a leading role in developing the new undergraduate major programme *Computational Data Science (CDAS)*, which was launched in Fall 2022. The programme is a joint programme between Department of Computer Science and Engineering and Department of Statistics at CUHK, featuring a strong inter-disciplinary curriculum with software development and data analytics trainings.

3.2 Student/Postdoc Supervision

3.2.1 Current Graduate Students

1. Ren, Yuanming, PhD 2023 - now (Graduated from TJU)
2. Liu, Qingxiu, PhD 2023 - now (Graduated from CAS)
3. Puyang, Huancheng, PhD 2022 - now (Graduated from HKUST)
4. Ren, Yanjing, PhD 2022 - now (Graduated from UESTC)
5. Zhao, Jia, PhD 2022 - now (Graduated from CUHK)
6. Sheng, Siyuan, PhD 2021 - now (Graduated from CAS)
7. Cheng, Keyun, PhD 2020 - now (Graduated from CHUK)
8. Li, Jinhong, PhD 2019 - now (Graduated from ZJU)

3.2.2 Current Post-docs

1. Chan, Hoi Wan (Helen), Sep 2019 - now (PhD from CUHK)

3.2.3 Graduated Students

1. Tang, Kaicheng, MPhil 2022
Thesis: "Balancing Repair Bandwidth and Sub-packetization in Erasure-Coded Storage via Elastic Transformation."
First employment: Alibaba Group.
2. Wang, Qiuping, PhD 2022
Thesis: "Enabling High Performance in Large-Scale Block Storage Systems through New Caching and Garbage Collection Algorithms."
First employment: Software engineer, Alibaba Group.
3. Yang, Zuoru, PhD 2022
Thesis: "Enhancing Secure Deduplication with Tunability and Shielded Execution for Outsourced Storage."
First employment: Senior research engineer, Huawei Technologies.
4. Ma, Pengcheng, MPhil 2022
Thesis: "A Performance Evaluation Study of Triple-Fault-Tolerant Erasure Codes."

5. Cheng, Zhinan, PhD 2022
Thesis: “Fault-tolerant Designs for Distributed Stream Processing.”
First employment: Senior Software Engineer, Tencent.
6. Han, Shujie, PhD 2021
Thesis: “Characterization and Prediction of Disk Drive Failures at Scale.”
First employment: Boya Postdoc Fellow, Peking University.
7. Wong, Tsz Fung (Ivan), MPhil 2021
Thesis: “Enabling Heterogeneity-Aware Raft Consensus for Distributed Key-Value Storage.”
First employment: Software Analyst, Ahsay Backup, Hong Kong.
8. Tang, Lu, PhD 2020
Thesis: “Invertible Sketches for Network Measurement at Scale.” First employment: Postdoc, CUHK.
Next (from March 2021): Assistant Professor, Xiamen University.
9. Li, Xiaolu, PhD 2020
Thesis: “Enabling Optimal Repair and Unified Management of Erasure Coding in Distributed Storage Systems.”
First employment: Postdoc, CUHK
Next (from February 2022): Lecturer, Huazhong University of Science and Technology.
10. Zhang, Mi, PhD 2019
Thesis: “On the Reliability Analysis and Robust Straggler Tolerance of Data-Center Storage with Erasure Coding.”
First employment: Postdoc, CUHK.
Next (from September 2021): Assistant Researcher, Institute of Computing Technology, Chinese Academy of Sciences.
11. Chan, Hoi Wan (Helen), PhD 2018
Thesis: “Log-structured Flash-based Storage with Endurance, Reliability, and Performance Guarantees”
First employment: R&D engineer, Unissoft Inc., Hong Kong
Current (from September 2019): Postdoc, CUHK
12. Qin, Chuan, PhD 2016
Thesis: “Unifying Security and Deduplication in Cloud Storage”
First employment: Researcher, Rakuten Institute of Technology, Tokyo, Japan
13. Yiu, Man Tung (Matt), MPhil 2016
Thesis: “Erasure Coding in In-Memory Key-Value Storage”
First employment: Stepcase, Hong Kong. (Stepcase is the company that built Lifehack)
14. Li, Runhui, PhD 2015
Thesis: “Enabling Efficient and Dependable Clustered File Systems through New Erasure Coding Techniques.”
First employment: Postdoc at CUHK
Current (from 2018): Storage expert, Sangfor Technologies, China
15. Huang, Qun, PhD 2015
Thesis: “Efficient Algorithms and Systems for High-Speed Network Traffic Analytics.”
First employment: Researcher, Future Network Theory Lab, Huawei Technologies, Hong Kong.
Current (from June 2020): Assistant Professor, Peking University, China
Award: HKPHD Fellowship Recipient
16. Xu, Min, MPhil 2015
Thesis: “Even Data Placement for Load Balance in Reliable Distributed Deduplication Storage.”
First employment: PhD student in Computer Science at University of Chicago
17. Chan, Chun Wing (Jeremy), MPhil 2014
Thesis: “Enabling High-Performance Erasure-Coded Storage Systems via Workload-Aware Designs.”
First employment: Analyst, Barclays Investment Bank, Hong Kong.

18. Ding, Qian, MPhil 2014
Thesis: "TWEEN: Toward Write Efficiency and Endurance in SSD RAID using Non-volatile Cache."
First employment: Software Engineer, Google, Mountain View, US.
19. Li, Yan Kit (Wilson), MPhil 2014
Thesis: "Performance and Security of Deduplication Storage Systems."
First employment: Research assistant at CUHK.
20. Lin, Jian, MPhil 2013
Thesis: "Augmenting Cloud File Systems with Erasure Coding: A Practical Perspective"
First employment: Epic, Wisconsin, US.
21. Tang, Chung Pan (Dick), MPhil 2013
Thesis: "A Tunable Version Control System for Virtual Machines in an Open-Source Cloud"
First employment: Harmonic Inc., Hong Kong
22. Chen, Chuk Hin (Henry), MPhil 2012
Thesis: "Practical Data Integrity Protection in Network-coded Cloud Storage."
First employment: Research assistant at CUHK.
Next: Ph.D. student in Computer Science at UT Austin (from Fall 2013)
23. Ng, Chun Ho, MPhil 2012
Thesis: "Live Deduplication Storage of Virtual Machine Images in an Open-Source Cloud."
First employment: Research assistant at CUHK.
Next: Epic, Wisconsin, US
24. Ma, Mingcao, MPhil 2012
Thesis: "Kernel-Space Inline Deduplication File Systems for Virtual Machine Image Storage."
First employment: NetDimensions, Hong Kong
25. Rahumed, Arthur, MPhil 2012
Thesis: "Data Security and Reliability in Cloud Backup Systems with Deduplication."
First employment: Harmonic Inc., Hong Kong
26. Choi, Fu-Wing (Harry), MPhil 2011
Thesis: "An Extensible Design of a Load-Aware Virtual Router Monitor in User Space."
First employment: Software engineer at Harmonic Inc.
27. Fung, Siu-Yuen (Ben), MPhil 2011
Thesis: "DeRef: A Privacy-Preserving Defense Mechanism Against Request Forgery Attacks"
First employment: Management trainee at Standard Chartered Bank.
Award: Best paper award in IEEE TrustCom-11.
28. Tang, Yang, MPhil 2011
Thesis: "FADE: Secure Overlay Cloud Storage with Access Control and File Assured Deletion."
First employment: Ph.D. student in Computer Science at Columbia University.
Award: Merit prize, Professor Charles K. Kao Student Creativity Awards, 2011

3.2.4 Post-doc Alumni

1. Wu, Si, Postdoc, January 2019 - July 2020
First employment: Associate Researcher, University of Science and Technology of China, China.
2. Shen, Zhirong, Postdoc, September 2017 - August 2019
First employment: Associate Professor, Xiamen University, China
3. Li, Jingwei, Postdoc, January 2015 - February 2016
First employment: Assistant Professor, University of Electronic Science and Technology of China (UESTC), China

4. Li, Mingqiang, Postdoc, March 2013 - February 2015
First employment: Advisory Researcher, Lenovo Group Limited
5. Hu, Yuchong, Postdoc, December 2013 - January 2015
First employment: Associate professor, Huazhong University of Science and Technology, China
6. Li, Yongkun, Postdoc (affiliated with Institute of Network Coding), September 2012 - December 2013
First employment: Associate researcher at University of Science and Technology of China

3.3 Classroom Teaching

3.3.1 Courses Taught

- **AIST3020 Introduction to Computer Systems:** Spring 2024, Spring 2023, Spring 2022, Spring 2021
 - An undergraduate course that covers the basics of computer systems.
- **CSCI4180 Introduction to Cloud Computing and Storage:** Fall 2023, Fall 2022, Fall 2021, Fall 2020, Fall 2019, Fall 2018, Fall 2017, Fall 2015, Fall 2014, Fall 2013, Fall 2012, Spring 2012 (1st half)
 - A senior undergraduate course that covers the applied cloud computing and storage methodologies and principles. It also covers real-life cases from enterprises (e.g., Google, Facebook, Dropbox, Amazon, Yahoo!, NetApp, etc.) to support the teaching materials.
- **CSCI5550 Advanced File and Storage Systems:** Fall 2023, Fall 2021, Spring 2018
 - A graduate course that covers the advanced topics in file and storage systems, such as file systems, SSDs, in-memory storage, RAID and erasure coding, etc.
- **ESTR4106 Introduction to Cloud Computing and Storage (Elite class):** Fall 2019, Fall 2018, Fall 2017
 - An elite version of CSCI4180. It covers state-of-the-art cloud computing platforms, such as Spark, Storm, and Kafka.
- **CSCI4430 Data Communications and Computer Networks:** Spring 2021, Spring 2020, Spring 2019, Spring 2018, Spring 2016, Spring 2015 (2nd half), Spring 2013 (2nd half), Spring 2012 (2nd half), Spring 2011 (1st half), Spring 2010
 - A senior undergraduate course that covers the computer networking methodologies and principles. It covers not only the classical layering design in computer networking, but also some applied techniques such as socket programming, multi-threaded programming, and Linux router designs.
- **ESTR4120 Data Communications and Computer Networks (Elite class):** Spring 2020, Spring 2019, Spring 2018
 - An elite version of CSCI4430. It covers several network security topics from both attack and defense perspectives.
- **ENGG5105/CSCI5470 Computer and Network Security:** Spring 2014, Spring 2013, Fall 2011, Fall 2010
 - A graduate course that covers the advanced topics in computer and network security.
- **CSCI4430S Data Communications and Computer Networks (Honor class):** Spring 2011 (2nd half)
 - An elite version for CSCI4430.
- **CMSC5709 Advanced Topics in Internet Technology:** Spring 2011 (1st half)
 - A graduate course that covers the literature in Internet Technology.
- **GEUC4011 Senior Seminar:** Fall 2013
 - A college education course for senior undergraduate students.

4 Contributions in Services

4.1 Internal Services

4.1.1 Department Services

- Vice Department Chairman (Undergraduate), 2022-24
- Programme Director, Digital Literacy Education (University Core Course), 2022-24
- Chair, Curriculum Committee, 2017-24
- Member, Executive Committee, 2017-24
- Chair, Dual-Degree Programme (DDP) Computer Science (CS) Management Committee, 2019-24
- Member, MSc Programme AI Taskforce, 2022-24
- Member, AIST Committee, 2018-19
- Member, Board of Information Engineering, 2018-19
- Member, Graduate Panel, 2009-11, 2014-16
- Chair, Undergraduate Academic Counseling, 2014-16
- Member, Staff-Student Consultative Committee, 2014-15
- Member, Curriculum Committee, 2012-14
- Member, MSc Programme Management Committee, 2011-12, 2015-16
- Member, Programmes for Secondary School Students, 2011-12
- Chairperson, Computer and Network, QoS and Security Committee, 2010-11
- Member, Academic Counseling and Course Selection for Mainland Students, 2010-11

4.1.2 Faculty Services

- Member, Faculty Curriculum Committee, 2017-24
- Member, Faculty Academic Advising (CDAS year 1 students) (Level 2), 2022-23
- Member, Faculty Committee on Academic Advising, 2014-15
- Academic advisor, Faculty Academic Advising (Level I), 2012-15
- Member, Subject Panel for ENGG1820 Engineering Internship, 2011-13
- Member, Teaching and Learning Committee for Engineering Design II, 2010-11

4.1.3 College Services

- Departmental Coordinator, United College, 2011-12, 2014-15, 2018-19, 2021-22, 2023-24
- Teacher, General Education GEUC4011 Senior Seminar, Fall 2013

4.1.4 University Services

- Member, Central Cluster Management Committee, 2018-present

4.1.5 Awards

- Department Exemplary Service Awards, 2017-18, 2020-21

4.2 External Services

- External Reviewer, School of Science and Engineering, The Chinese University of Hong Kong (Shenzhen), China, 2017-present