

Curriculum Vitae

Kai Li

Department of Computer Science
35 Olden Street
Princeton University
Princeton, N.J. 08544-2087
(609) 258-4637
E-mail: li@cs.princeton.edu

Education

- Ph.D. 1986, Yale University. Yale Fellowships.
Dissertation: *Shared Virtual Memory on Loosely Coupled Multiprocessors*.
Advisers: Paul Hudak and Alan Perlis.
- M.S. 1983, M.Phil. 1984, Yale University.
- M.S. 1981, University of Science and Technology of China, Peking, P.R.C.
- B.S. 1977, Jilin University, Changchun, P.R.C. Graduate with high honors.

Research Interests

Parallel and distributed systems, operating systems, storage systems, computer architecture, large-scale data analysis, and knowledge base construction.

Full-Time Employment

- Paul M. Wythes '55, P'86 and Marcia R. Wythes P'86 Professor, Department of Computer Science, Princeton University, since 2009.
- Charles Fitzmorris Professor, Department of Computer Science, Princeton University, 2002-2009.
- Chief Scientist, Data Domain, Inc. 2002.
- Chief Executive Officer, Data Domain, Inc. 2001-2002.
- Professor, Department of Computer Science, Princeton University, 1995-2001.
- Associate professor, Department of Computer Science, Princeton University, 1992-1995.
- Assistant professor, Department of Computer Science, Princeton University, 1986-1992.

Part-Time and Consulting

- Vice President, Asian American Scholar Forum (nonprofit), since 2021.
- EMC Visiting Professor, Tsinghua University, 2008.
- Honorary visiting professor, Tsinghua University, 2005-2008.
- Honorary visiting professor, Institute of Computing Technology, Chinese Academy of Sciences, since 2005.
- Chief Scientist, Data Domain, Inc. 2003-2009.
- Entrepreneur-In-Resident, New Enterprise Associates, fall 2001.
- Honorary Visiting Professor, Institute of Software, Academy of Sciences of China, 2001.
- Visiting Professor, Department of Computer Science, Stanford University, 2000.
- Visiting Professor, Department of Computer Science, Stanford University, 1996.
- Consultant, Intel Microcomputer Research Labs, since 1996.
- Consultant, AT&T Bell Laboratories, summer 1995.
- Consultant, NEC Research Institute, since 1994.
- Consultant, Matsushita Information Technology Laboratory, 1991-94.
- Consultant, DEC Systems Research Center, 1989.
- Consultant, Intel Supercomputer Systems Division, 1987.

- Visiting assistant professor, Department of Electrical Engineering and Department of Computer Science, University of Toronto, Fall 1988.
- Visiting scientist, DEC Systems Research Center, summer 1987 and 1988.

Test-of-Time / Most Influential Paper Awards

- *Test-of-time award in 50 years of International Conference on Parallel Processing (ICPP)*, for paper “IVY: A Shared Virtual Memory System for Parallel Computing”(published in 1988), 2021.
- *Test-of-time award of USENIX File and Storage Technologies (FAST)*, for paper “Avoiding the Disk Bottleneck in the Data Domain Deduplication File System” (published in 2008), 2020.
- *Longuet-Higgins prize, 10-year best paper of Computer Vision and Pattern Processing (CVPR)* for paper “ImageNet: a large-scale hierarchical image database” (published in 2009), 2019.
- *Test-of-Time Award of Very Large Data Base (VLDB)*, for paper “Multi-Probe LSH: Efficient Indexing for High-Dimensional Similarity Search” (published in 2007). 2017.
- *ACM SIGOPS Hall of Fame Award*, for paper “Memory Coherence in Shared Virtual Memory Systems” (published in ACM TOCS in 1989), 2012.
- *50 most-influential papers in 20 Years of ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, for paper “Real-time Concurrent Collection on Stock Multiprocessors” (published in 1988), 2004.
- *43 most-influential papers in 25 Years of ACM/IEEE International Symposium of Computer Architecture (ISCA)*, for paper “Virtual Memory Mapped Network Interface for the Shrimp Multicomputer” (published in 1994). 1998.

Honors

- Foreign Member of Chinese Academy of Engineering, 2017.
- Member of National Academy of Engineering, 2012.
- Member of Washington State Academy of Sciences, 2012.
- IEEE Fellow, 2011.
- ACM Fellow, 1998.

Awards

- Overseas Outstanding Contribution Award, China Computer Federation, 2008.
- Commendation List for Outstanding Teaching, School of Engineering and Applied Science, Princeton University, 2006.
- IBM Partnership Award, 1998 and 2000.
- Distinguished Achievement Award, Chinese Institute of Engineers, 2012.
- Best paper award, ACM Measurement and Modeling of Computer Systems, 1995.
- Excellence in Teaching Award, Undergraduate Engineering Council, Princeton University, 1994.

Boards

- Director, Board of Directors, Asian American Scholar Forum, since 2021.
- Member, Engineering Leadership Council, School of Engineering and Applied Science, Yale University, since 2021.
- Advisory Council Member, Computer Science, Tsinghua University, 2017-2021.
- International Review Panel, Electrical Engineering and Computer Science, Peking University, 2017.

- Director, Board of Directors, Moqi, Inc., 2017-2021.
- Technical advisory board member, DELL/EMC, since 2017-2018.
- Advisory Board Chair, Jilin University, China, 2015.
- Advisory Board Member, Open Innovation Center, Samsung Semiconductor, Inc. 2012-2016.
- Advisory board member, The Technion Computing Engineering Center, Technion Israel Institute of Technology, 2011-2012.
- Technical advisory board member, Inphi Corporation, since 2010-2011.
- Technical advisory board member, EMC Corporation, 2009-2015.
- Advisory board chair, Intel Beijing Research Lab, 2009-2014.
- Director, Board of Directors, Pattern Insight, Inc. 2007-2012.
- Advisory board member, School of Computer and Communication Sciences (I&C), Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland, 2004-2018.
- Advisory board member, Computer Science Department, Hong Kong University of Science and Technology, Hong Kong, 2004-2007.
- Advisory board member, TeraGrid Technical Advisory Board, Argonne National Labs, NCSA and UCSD Supercomputing Center, 2002-2004.
- Director, Board of Directors, Data Domain, Inc., 2001-2009.
- Advisory board member, Board of Advisory council, Intel Microcomputer Technology Labs, 2000-2004.
- Advisory board member, DOE Plasma Sciences Advanced Computing Institute (PSACI), 1999-2009.

Professional Services

- Member, Qiu-Shi Science Prize, 2017.
- Member, Science Committee, Future Science Prize, since 2015
- Chair, ACM SIGOPS Mark Weiser Award committee, 2010.
- Member, ACM SIGOPS Mark Weiser Award committee, 2008-2009.
- Panel Member, Research Assessment Exercise (RAE), University Grants Committee, Hong Kong, 2006.
- Member, Committee of Visitors, National Science Foundation, 2005.
- Editorial board member, *International Journal of Parallel Programming*, since 1993.
- Associate Editor, *IEEE Transactions on Parallel and Distributed Systems*, 1994-97.
- Guest Editor, Special Issue on "Large-Format Displays," *IEEE Computer Graphics & Applications*. July 2000.
- Editorial board member, *Journal of Software*, since 2001.
- Editorial board member, *International Journal of Security and Networks*, since 2005.
- Program committee member, USENIX 13th Conference on File and Storage Technologies, 2013, 2016.
- Program committee member, ACM International Conference on Multimedia Retrieval, 2012.
- Program committee member, ACM/IEEE International Symposium of Computer Architecture, 2012.
- Program committee member, ACM Workshop of Management of Big Data Systems, 2012.
- Program committee member, ACM/IEEE International Symposium of Computer Architecture, 2011.
- Program committee member, ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems, 2010.

- Program committee member, ACM SIGMM the 11th International Conference on Multimedia Information Retrieval, 2010.
- Co-general chair, ACM/IEEE International Symposium of Computer Architecture, 2008.
- Program committee member, IEEE International Symposium on High Performance Computer Architecture, 2008.
- Program committee member, ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems, 2007.
- Program committee member, ACM SIGOS Eurosys Conference, 2007.
- Program committee member, ACM Workshop on Continuous Archival and Retrieval of Personal Experiences, 2006.
- Program committee member, IEEE International Symposium on High Performance Computer Architecture, 2006.
- Program committee member, IEEE International Parallel and Distributed Processing Symposium, 2006.
- Program committee member, ACM International Conference on Computing Frontiers, 2006.
- Co-general chair, IEEE High Performance Computing Asia, 2005.
- Program committee member, USENIX Conference on File And Storage Technology (FAST), 2005.
- Program committee member, ACM Workshop on Continuous Archival and Retrieval of Personal Experiences, 2005.
- Program committee member, IEEE Workshop on Remote Direct Memory Access, 2005.
- Co-program committee chair, ACM/IEEE Symposium on Architecture for Networking and Communications Systems, 2005.
- Program committee member, ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems, 2005.
- Steering committee member, ACM/IEEE International Symposium of Computer Architecture, 2004.
- Program committee member, ACM Workshop on Continuous Archival and Retrieval of Personal Experiences, 2004.
- Program committee member, IEEE Workshop on Remote Direct Memory Access, 2004.
- Program committee chair, ACM/IEEE International Symposium of Computer Architecture, 2003.
- Program committee vice chair, ACM Supercomputing conference, 2002
- Program committee member, ACM International Conference on Architectural Support for Operating Systems and Programming Languages, 2000.
- Program committee member, ACM Supercomputing Conference, 2000.
- Program committee member, ACM International Conference on Supercomputing, 2000.
- Program committee chair, IEEE High-Performance Computer Architecture, 2000.
- Program committee member, ACM/IEEE International Symposium of Computer Architecture, 1999.
- Program committee member, ACM International Conference on Supercomputing, 1999.
- Program committee member, ACM SIGOS/USENIX Symposium on Operating Systems Design and Implementation, 1999.
- Program committee member, ACM Symposium on Principles of Distributed Computing, 1998.
- Program committee member, ACM Symposium on Parallel Algorithms and Architectures, 1998.
- Program committee Vice-Chair, International Conference on Parallel Processing, 1998.
- Program committee member, ACM SIGOS Symposium on Principles of Operating Systems, 1997.

- Program committee member, ACM Supercomputing conference, 1997.
- Program committee member, IEEE Hot Interconnects V, 1997.
- Program committee member, ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems, 1997.
- Program committee Vice-Chair, IEEE Symposium on Parallel and Distributed Processing. 1997.
- Program committee member, IEEE International Conference on Massively Parallel Processing using Optical Interconnections. 1996.
- Program committee Co-Chair, IEEE Hot Interconnects IV, 1996.
- Program committee member, IEEE Symposium on Parallel and Distributed Processing. 1996.
- Program committee member, IEEE Symposium on High Performance Computer Architecture, 1996.
- Program committee member, IEEE Hot Interconnects III, 1995.
- Program committee member, ACM Symposium on Parallel Algorithms and Architectures, 1995.
- Program committee member, ACM SIGOS Workshop on Hot Topics of Operating Systems, 1995.
- Program committee member, International Conference of Parallel Processing, 1995.
- Program committee member, IEEE International Parallel Processing Symposium, 1994.
- Program committee member, ACM Symposium on Principles of Distributed Computing, 1993.
- Program committee member, ACM/IEEE International Symposium on Computer Architecture, 1993.
- Program committee member, IEEE International Conference on Distributed Computing Systems, 1993.
- Program committee member, ACM SIGOS Symposium on Principles of Operating Systems, 1993.
- Program committee member, ACM/IEEE International Symposium on Computer Architecture, 1992.

Ph.D. graduates

1. Mark Greenstreet. Thesis: STARI: A Technique for High-Bandwidth Communication, 1993.
2. James Plank. Thesis: Efficient Checkpointing on MIMD Architectures, 1993.
3. Karin Petersen-Flynn. Thesis: Operating Systems Support for Modern Memory Hierarchies, 1993.
4. Pei Cao. Thesis: Application-Controlled File Caching and Prefetching, 1996.
5. Matthias Blumrich. Thesis: Network Interface for Protected, User-Level Communication, 1996.
6. Stefanos N. Damianakis. Thesis: Efficient Connection-Oriented Communication on High-Performance Networks. 1998.
7. Liviu Iftode. Thesis: Home-Based Shared Virtual Memory, 1998.
8. Bin Wei. Thesis: A Multiple-Port Distributed Frame Buffer on Multicomputers, 1998.
9. Yuanyuan Zhou. Thesis: Memory Management for Networked Servers, 2000.
10. Yuqun Chen. Thesis: Building A Scalable High-Resolution Display Wall, 2000.
11. Sanjeev Kumar. Thesis: ESP: A Language for Programmable Devices, 2002.
12. Han Chen. Thesis: Scalable and Ultra-High Resolution MPEG Video Delivery on Tiled Displays. 2003.
13. Qin (Christine) Lv. Thesis: Similarity Search for Large-Scale Image Datasets. 2006.
14. Zhiyan (Cathy) Liu. Thesis: Scalable Isosurface Visualization. May 2008.
15. Matthew Hibbs (co-advised with Prof. Olga Troyanskaya). Thesis: Analysis and Visualization of Large-Scale Gene Expression Microarray Compendia, 2007.
16. Junwen Lai (co-advised with Dr. Randy Wang). Thesis: Virtualizing Network File Systems, 2008.
17. Christian Bienia. Thesis: Benchmarking Modern Multiprocessors, 2010.
18. William Josephson. Thesis: A Direct-Access File System for a New Generation of Flash Memory, 2011.
19. Wei Dong. Thesis: High-Dimensional Similarity Search for Large Datasets, 2011.
20. Zhe Wang. Thesis: Similarity Search with Multimodal Data, 2011.
21. Jia Deng (co-advised with Prof. Fei-fei Li). Thesis: Large Scale Visual Recognition, 2012.
22. Yida Wang. Thesis: Large-scale analyses of functional interactions in the human brain. 2016.
23. Qian Zhu (co-advised with Prof. Olga Troyanskaya).
Thesis: Detecting Gene Similarities using Large-Scale Content-Based Search Systems. 2016.
24. Linpeng Tang. Thesis: Efficient Processing and Delivery of Multimedia Data. 2018.
25. Lifang Cheng (co-advised with Prof. Barbara Engelhardt). Thesis: Large-Scale Multi-Output Gaussian Processes for Clinical Decision Support. 2019.
26. Sachin Ravi (co-advised with Prof. Jonathan Cohen). Thesis: Meta Learning for Data and Processing Efficiency. 2019.

Recent Keynote/Distinguished Lectures

- Distinguished lecture, “Machine Learning Meets Systems.” David Cheriton School of Computer Science, University of Waterloo, May 2022.
- Distinguished lecture, “Disruptive Research and Innovation.” The 14th Annual Conference of Chinese Academy of Engineering. Beijing, China. May 2018.

- Keynote (with Jonathan Cohen), “Understanding the Human Brain: A Computational Challenge of Truly Cosmic Proportions,” Intel HPC Developer Conference. November 2016.
- Distinguished lecture, “Disruptive Innovation: Data Domain Experience,” Jonsson School's 30th Anniversary Distinguished Lecture Series. University of Texas at Dallas. October 2016.
- Keynote, “Parallel Computing for Analyzing Functional Interactions of the Human Brain,” 11th Conference on Advanced Computer Architecture. August 2016.
- Keynote, “Disruptive Research and Innovation,” 30th IEEE International Parallel and Distributed Processing Symposium (IPDPS). May 2016.
- Keynote, “A Real-Time, Closed Loop fMRI Data Analysis System for Neuroscience Experiments,” Intel HPC Developer Conference, November 2015.
- Keynote, “Deduplication Storage Systems: Reducing Space and Energy,” Frontiers in Energy-Efficient Computing and Applications. September 2015.
- Distinguished lecture, “Disruptive Research and Innovation,” Future Forum, Beijing China, June 2015.
- Distinguished lecture, “How to Succeed in Research and Innovation,” School of Computer Science and Technologies, Jilin University, China, September 2014.
- Distinguished lecture, “Disruptive Innovation: Data Domain Experience,” Department of Computer Science, University of Illinois at Urbana-Champaign, October 2013.
- Distinguished Lecture, “Dual of Research and Innovation,” Department of Computer Science, Ohio State University, April 2012.
- Keynote, “Dual of Research and Innovation,” Forum of Computer Science Department Heads, Wuhan, August 2011.
- Keynote, “Challenges in Building A Deduplication Storage System,” The 2nd ACM SIGOPS Asia-Pacific Workshop on Systems (APSys 2011), Shanghai, China, July 2011.
- Keynote, “Challenges in Building A Commercial Deduplication Storage System,” The 4th Annual International Systems and Storage Conference, Haifa, Israel, May 2011.
- Distinguished lecture, “Building a Commercial Deduplication Storage System,” Department of Computer Science, Purdue University, November 2010.
- Distinguished lecture, “Building a Commercial Deduplication Storage System,” Microsoft Research Silicon Valley, July 2010.
- Distinguished lecture, “Deduplication Storage System,” Department of Computer Science, Wayne State University, March 2009.
- Keynote, “Deduplication Storage System,” Annual Conference of China Computer Federation, June 2008.
- Distinguished lecture, “A Tapeless World Can Be Inexpensive,” School of Computer Science, University of Illinois at Urbana Champaign, April 2007.
- Keynote speech, “Next-Generation Data Protection: Progress and Challenges,” 50th Anniversary of Institute of Computing Technology, Chinese Academy of Sciences, China, October 2006.
- Distinguished lecture, “Content-Based Similarity Search,” Tsinghua University, China, December 2006.
- Distinguished lecture, “Content-Based Similarity Search,” Institute of Computing Technology, Chinese Academy of Sciences, China, October 2005.
- Distinguished lecture, “Princeton Scalable Display Wall,” Zhongshan University, China, January 2005.
- Distinguished lecture, “Building and Using A Scalable Display Wall,” University of Oslo, September 2001.

- Distinguished lecture, "Building and Using A Scalable Display Wall," University of Tromso, September 2001.
- Keynote speech, "Challenges and Early Experiences in Building and Using A Scalable Display Wall," IEEE Parallel Visualization and Graphics Symposium, October 1999.
- Distinguished Lecture, "Inexpensive Visual Computing," Department of Computer Science, University of Illinois at Urbana-Champaign, October 1999.
- Distinguished lecture, "Inexpensive Visual Computing," Department of Electrical Engineering, University of Toronto, September 1999.
- Distinguished lecture, "Inexpensive Visual Computing," Department of Computer Science, University of Chicago, March 1999.
- Distinguished lecture, "Inexpensive Visual Computing," Department of Computer Science, University of Tennessee, February 1999.

Patents

1. Efficient computation of sketches, (with Ming Benjamin Zhu). U.S. Patent 9,823,975. 2017.
2. Stream locality delta compression, (with Mark Huang, Philip Shilane, Grant Wallace, Nitin Garg, Edward K. Lee, Benjamin Ming Zhu), U.S. Patent 9,690,802. 2017.
3. Efficient Data Storage System (with Ming Benjamin Zhu and Hugo Patterson), U.S. Patent 9,454,318. 2016.
4. Efficient Data Storage System (with Ming Benjamin Zhu and Hugo Patterson), U.S. Patent 9,081,728. 2015.
5. Efficient Data Storage System (with Ming Benjamin Zhu and Hugo Patterson), U.S. Patent 8,612,696. 2013.
6. Cluster storage using delta compression (with Hugo Patterson, Ming Benjamin Zhu, Sazzala Venkata Reddy, Umesh Maheshwari, Edward K. Lee). U.S. Patent 8,312,006. 2012.
7. Efficient Data Storage System (with Ming Benjamin Zhu and Hugo Patterson), U.S. Patent 8,275,955. 2012.
8. Partitioning a data stream using embedded anchors (with Umesh Maheshwari and Hugo Patterson). U.S. Patent 8,234,413. 2012
9. Partitioning a data stream using embedded anchors (with Umesh Maheshwari and Hugo Patterson). U.S. Patent 7,979,584. July 2011.
10. Similarity search system with compact data structures (with Qin Lv and Moses Charikar). U.S. Patent 7,966,327. June 2011.
11. Cluster storage using delta compression (with Hugo Patterson, Ming Benjamin Zhu, Sazzala Venkata Reddy, Umesh Maheshwari, Edward K. Lee). U.S. Patent 7,962,520. June 2011.
12. Efficient computation of sketches, (with Ming Benjamin Zhu). U.S. Patent 7,844,652. 2010.
13. Efficient Data Storage System (with Ming Benjamin Zhu and Hugo Patterson). U.S. Patent 7,769,967. August 2010.
14. Network file system-based data storage system (with Hugo Patterson; Zhu; Ming Benjamin, Allan Bricker. Richard Johnsson, Sazzala Reddy, Jeffery Zabarsky). U.S. Patent 7,747,581. June 2010.
15. Network file system-based data storage system (with Hugo Patterson; Zhu; Ming Benjamin, Allan Bricker. Richard Johnsson, Sazzala Reddy, Jeffery Zabarsky). U.S. Patent 7,689,633. March 2010.
16. Efficient data storage using resemblance of data segments (with M. Zhu, U. Maheshwari, Z. Yang). U.S. Patent 7,562,186. 2009.
17. Efficient Data Storage System (with Ming Benjamin Zhu and Hugo Patterson). U.S. Patent 7,434,015. October 2008.
18. Efficient Data Storage System (with Ming Benjamin Zhu and Hugo Patterson). U.S. Patent 7,373,464. May 2008.
19. Efficient Data Storage System (with Ming Benjamin Zhu and Hugo Patterson). U.S. Patent 7,305,532. December 2007.
20. Efficient Data Storage System (with Ming Benjamin Zhu and Hugo Patterson). U.S. Patent 7,065,619. June 2006.
21. Archival Data Storage System and Method (with Howard Lee). U.S. Patent 7,007,141. 2006.
22. Efficient Data Storage System (with Ben Zhu and Hugo Patterson). U.S. Patent 6,928,526. August 2005.

23. Optical blending for multi-projector display wall systems (with Yuqun Chen and Timothy Housel). U.S. Patent 6,570,623. 2000.
24. Method for Improving Cache Locality of A Computer Program (with Jan Edler and James Philbin). U.S. Patent 5,724,586. 1996.
25. User-Level Direct Memory Access (with Matthias Blumrich, Cezary Dubnicki, and Edward Felten). U.S. Patent 5,659,798. 1996.
26. Real Time, Concurrent Garbage Collection System and Method (With Andrew Appel and John Ellis). U.S. Patent 5,088,036. 1989.

Publications

27. Yangsibo Huang, Chun-Yin Huang, Xiaoxiao Li, Kai Li. A Dataset Auditing Method for Collaboratively Trained Machine Learning Models. *Special Issues of IEEE Transaction on Medical Imaging*. To appear.
28. Yangsibo Huang, Samyak Gupta, Zhao Song, Kai Li, Sanjeev Arora. Evaluating Gradient Inversion Attacks and Defenses. Book chapter, *Federated Learning: Theory and Practice*. To appear.
29. Samyak Gupta, Yangsibo Huang, Zexuan Zhong, Tianyu Gao, Kai Li, Danqi Chen. Recovering Private Text in Federated Learning of Language Models. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
30. Yangsibo Huang, Samyak Gupta, Zhao Song, Kai Li, and Sanjeev Arora. Evaluating Gradient Inversion Attacks and Defenses in Federated Learning. In *Advances in Neural Information Processing Systems (NeurIPS)* 2021.
31. Yangsibo Huang, Xiaoxiao Li, and Kai Li. EMA: Auditing Data Removal from Trained Models. In *International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI)* 2021.
32. Manoj Kumar, Michael J. Anderson, James W. Antony, Christopher Baldassano, Paula P. Brooks, Ming Bo Cai, Po-Hsuan Cameron Chen, Cameron T. Ellis, Gregory Henselman-Petrusek, David Huberdeau, J. Benjamin Hutchinson, Y. Peeta Li, Qihong Lu, Jeremy R. Manning, Anne C. Mennen, Samuel A. Nastase, Hugo Richard, Anna C. Schapiro, Nicolas W. Schuck, Michael Shvartsman, Narayanan Sundaram, Daniel Suo, Javier S. Turek, David Turner, Vy A. Vo, Grant Wallace, Yida Wang, Jamal A. Williams, Hejia Zhang, Xia Zhu, Mihai Capota, Jonathan D. Cohen, Uri Hasson, Kai Li, Peter J. Ramadge, Nicholas B. Turk-Browne, Theodore L. Willke, and Kenneth A. Norman. BrainIAK: The Brain Imaging Analysis Kit. *Aperture Neuro*. 2021.
33. Zhenyu Song, Daniel S. Berger, Kai Li, Wyatt Lloyd. Learning Relaxed Belady for Content Distribution Network Caching. *USENIX Conference on Network Systems Design and Implementation*. 2020.
34. Yangsibo Huang, Zhao Song, Danqi Chen, Kai Li, and Sanjeev Arora. TextHide: Tackling Data Privacy in Language Understanding Tasks. In *Conference on Empirical Methods in Natural Language Processing (Findings of EMNLP)* 2020.
35. Yangsibo Huang, Zhao Song, Danqi Chen, Kai Li, and Sanjeev Arora. InstaHide: Instance-hiding Schemes for Private Distributed Learning. In *International Conference on Machine Learning (ICML)* 2020.
36. Li-Fang Cheng, Bianca Dumitrascu, Gregory Darnell, Corey Chivers, Michael Draugelis, Kai Li and Barbara Engelhardt . Sparse Multi-Output Gaussian Processes for Medical Time Series Prediction. In *BMC Medical Informatics and Decision Making*, 2020.
37. Zhen Jia, Aleksandar Zlateski, Fredo Durand, Kai Li. The Anatomy of Efficient FFT and Winograd Convolutions on Modern CPUs. *International Conference on Supercomputing (ICS)*. 2019.

38. Li-Fang Cheng, Bianca Dumitrascu, Michael Zhang, Corey Chivers, Michael Draugelis, Kai Li and Barbara Engelhardt. Patient-Specific Effects of Medication using Latent Force Models with Gaussian Processes. In *the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2020.
39. Zhen Jia, Aleksandar Zlateski, Fredo Durand, Kai Li. Towards Optimal Winograd Convolution on Manycores. SysML Conference. 2018.
40. Aleksandar Zlateski, Zhen Jia, Kai Li, Fredo Durand. A Deeper Look at FFT and Winograd Convolutions. SysML Conference. 2018
41. Zhen Jia, Aleksandar Zlateski, Frédo Durand and Kai Li. Optimizing N-Dimensional, Winograd-Based Convolution for Manycore CPUs. Proceedings of ACM Principles and Practice of Parallel Programming, 2018.
42. Li-Fang Cheng, Gregory Darnell, Bianca Dumitrascu, Corey Chivers, Michael Draugelis, Kai Li and Barbara Engelhardt. Sparse Multi-Output Gaussian Processes for Medical Time Series Prediction. arXiv:1703.09112.
43. Qin Lv, William Josephson, Zhe Wang, Moses Charikar, Kai Li. Intelligent Probing for Locality Sensitive Hashing: Multi-Probe LSH and Beyond. Proceedings of the VLDB Endowment (10)12: 2021-2024. 2017.
44. Jonathan D. Cohen, Nathaniel Daw. Barbara Engelhardt, Uri Hasson, Kai Li, Yael Niv, Kenneth A. Norman, Jonathan Pillow, Peter Ramadge , Nicholas B. Turk-Browne, & Theodore L. Willke. Computational Approaches to fMRI Analysis. *Nature Neuroscience* (20): 304–313, 2017. doi:10.1038/nn.4499
45. Linpeng Tang, Qi Huang, Amit Puntambekar, Ymir Vigfusson, Wyatt Lloyd, Kai Li. Popularity Prediction of Facebook Videos for Higher Quality Streaming. USENIX Annual Technology Conference. 2017.
46. Vladimir Feinberg, Li-Fang Cheng, Kai Li and Barbara Engelhardt. Large Linear Multi-Output Gaussian Process Learning. arXiv:1705.10813.
47. Yida Wang, Bryn Keller, Mihai Capota, Michael J. Anderson, Narayanan Sundaram, Jonathan D. Cohen, Kai Li, Nicholas B. Turk-Browne, Theodore L. Willke. Real-time Full Correlation Matrix Analysis of fMRI Data. IEEE Conference on Big Data. 2016.
48. Yue Cheng, Fred Douglass, Philip Shilane, Michael Trachtman, Grant Wallace, Peter Desnoyers, Kai Li. Erasing Belady’s Limitations: In Search of Flash Cache Offline Optimality. In Proceedings of USENIX Annual Technical Conference, 2016.
49. Li-Fang Cheng, Gregory Darnell, Corey Chivers, Michael Draugelis, Kai Li and Barbara Engelhardt. Sparse Multi-Output Gaussian Processes for Medical Time Series Prediction in Machine Learning for Health Workshop, Neural Information Processing Systems (NIPS), 2016.
50. Qian Zhu, Aaron Wong, Arjun Krishnan, Miriam Aure, Alicja Tadych, Ran Zhang, David Corney, Casey Greene, Lars Bongo, Vessela Kristensen, Moses Charikar, Kai Li, Olga Troyanskaya. Targeted exploration and analysis of large cross-platform human transcriptomic compendia. *Nature Methods*. 12(3): 211-214, 2015.
51. Yida Wang, Michael J. Anderson, Jonathan D. Cohen, Alexander Heinecke, Kai Li, Nadathur Satish, Narayanan Sundaram, Nick B. Turk-Browne, and Theodore L. Willke, “Full Correlation Matrix Analysis of fMRI Data on Intel® Xeon Phi™ Coprocessors,” in SC’15 International Conference for High Performance Computing, Networking, Storage and Analysis, 2015.
52. Yida Wang, Jonathan D. Cohen, Nick Turk-Browne. Full correlation matrix analysis (FCMA): An unbiased method for task-related functional connectivity. *Journal of Neuroscience Methods*. 251: 108-119, 2015.
53. Linpeng Tang, Qi Huang, Wyatt Lloyd, Sanjeev Kumar, and Kai Li. RIPQ: Advanced Photo Caching on Flash for Facebook. In Proceedings of the 13th USENIX conference on File and Storage Technologies (FAST 15), February 2015.

54. Li-Fang Cheng, Michael Draugelis, Kai Li and Barbara Engelhardt. A Temporal-Difference Approach with Unsupervised Data Imputation for Early Sepsis Prediction. 2nd Workshop on Machine Learning for Clinical Data Analysis, Healthcare and Genomics (NIPS 2014)
55. Wei Dong, Kai Li, Zhe Wang. High Confidence Near-Duplicate Image Detection. ACM International Conference on Multimedia Retrieval (ICMR), 2012.
56. Gala Yadgar, Michael Factor, Kai Li and Assaf Schuster. *Management of Multilevel, Multiclient Cache Hierarchies with Application Hints*. ACM Transactions on Computer Systems (TOCS). Volume 29 Issue 2, May 2011.
57. Wei Dong, Moses Charikar, Kai Li. Efficient K-Nearest Neighbor Graph Construction for Generic Similarity Measures. In Proceedings of the 20th International World Wide Web Conference. Hyderabad, India. March 2011.
58. Wei Dong, Fred Douglis, Kai Li, Hugo Patterson, Sazzala Reddy, Philip Shilane. Tradeoffs in Scalable Data Routing for Deduplication Clusters. In Proceedings of the 9th USENIX Conference on File and Storage Technologies, 2011.
59. Christian Bienia and Kai Li. Fidelity and Scaling of the PARSEC Benchmark Inputs. In Proceedings of the 2010 IEEE International Symposium on Workload Characterization, December 2010.
60. Jia Deng, Alex Berg, Kai Li and Li Fei-Fei. What does classifying more than 10,000 image categories tell us? European Conference on Computer Vision, September 2010.
61. Christian Bienia and Kai Li. Characteristics of Workloads Using the Pipeline Programming Model. In Proceedings of the Third Workshop on Emerging Applications and Many-Core Architecture. Held in conjunction with the 37th International Symposium on Computer Architecture, June 2010.
62. William K. Josephson, Lars A. Bongo, David Flynn, and Kai Li. DFS A File System for Virtualized Flash Storage. 8th USENIX Conference on File and Storage Technologies, February 2010.
63. J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li and L. Fei-Fei, ImageNet: A Large-Scale Hierarchical Image Database. IEEE Conference on Computer Vision and Pattern Recognition (CVPR), June 2009.
64. J. Deng, K. Li, M. Do, H. Su, L. Fei-Fei, Construction and Analysis of a Large-Scale Image Ontology. Annual Conference of Vision Sciences Society (VSS), May 2009.
65. Hibbs, M.A.; Myers, C.L.; Huttenhower, C.; Hess, D.C.; Li, K.; Caudy, A.A.; Troyanskaya, O.G. Directing Experimental Biology: A Case Study in Mitochondrial Biogenesis, PLoS Computational Biology, 2009
66. Christian Bienia and Kai Li. PARSEC 2.0: A New Benchmark Suite for Chip-Multiprocessors. In Proceedings of the 5th Annual Workshop on Modeling, Benchmarking and Simulation, June 2009.
67. Christian Bienia, Sanjeev Kumar, Jaswinder Pal Singh and Kai Li. The PARSEC Benchmark Suite: Characterization and Architectural Implications. In Proceedings of the 17th International Conference on Parallel Architectures and Compilation Techniques. October 2008.
68. Wei Dong, Zhe Wang, Moses Charikar, Kai Li. Efficiently Matching Sets of Features with Random Histograms. In *Proceedings of the 16th ACM International Conference on Multimedia*. Vancouver, Canada. October 2008.
69. Wei Dong, Zhe Wang, William Josephson, Moses Charikar, Kai Li. Modeling LSH for Performance Tuning.. In *In Proceedings of ACM 17th Conference on Information and Knowledge Management (CIKM)*. Napa Valley, CA, USA. October 2008.
70. Brendan Collins, Jia Deng, Kai Li and Fei-Fei Li Towards scalable dataset construction: An active learning approach. In *Proceedings of the European Conference on Computer Vision (ECCV) 2008*.

71. Wei Dong, Moses Charikar, Kai Li. Asymmetric Distance Estimation with Sketches for Similarity Search in High-Dimensional Spaces. *Proceedings of the 31st Annual International ACM SIGIR Conference on Research & Development on Information Retrieval*. Singapore. July 2008.
72. Gala Yadgar, Michael Factor, Kai Li, and Assaf Schuster. MC2: Multiple Clients on a Multilevel Cache. *Proceedings of the 28th International Conference on Distributed Computing Systems (ICDCS)*. June 2008.
73. Benjamin Zhu, Kai Li and Hugo Patterson. Avoiding the Disk Bottleneck in a Deduplication Storage System. In *Proceedings of The 6th USENIX Conference on File and Storage Technologies (FAST'08)*. February 2008.
74. Matthew A. Hibbs, David C. Hess, Chad L. Myers, Curtis Huttenhower, Kai Li, and Olga Troyanskaya. Exploring the Functional Landscape of Gene Expression: Directed Search of Large Microarray Compendia. *Bioinformatics* 23(20):2692-2699; doi:10.1093/bioinformatics/btm403. October 2007.
75. Qin Lv, William Josephson, Zhe Wang, Moses Charikar, Kai Li. Multi-Probe LSH: Efficient Indexing for High-Dimensional Similarity Search. In *Proceedings of the 33rd International Conference on Very Large Data Bases (VLDB)*. Vienna, Austria. September 2007. **(10-year best paper award, VLDB, 2017)**
76. William Josephson, Ruby Lee and Kai Li. ISA Support for Fingerprinting and Eraser Codes. In *Proceedings of IEEE Application-Specific Systems, Architectures and Processors*. August 2007.
77. Zhe Wang, William Josephson, Qin Lv, Moses Charikar, Kai Li. Filtering Image Spam with Near-Duplicate Detection. In *Proceedings of the 4th Conference on Email and Anti-Spam (CEAS)*. Mountain View, CA, USA. August 2007.
78. Matthew Hibbs, Grant Wallace, Maitreya Dunham, Kai Li, and Olga Troyanskaya. Viewing the Larger Context of Genomic Data through Horizontal Integration. *Proceedings of the 11th IEEE International Conference of Information Visualization*, July 2007.
79. Grant Wallace and Kai Li. Virtually Shared Displays and User Input Devices. In *Proceedings of the 2007 USENIX Annual Technical Conference*. Pp 375-380. June 2007.
80. D.A. Batchelor, M. Beck, A. Becoulet, R.V. Budny, C.S. Chang, P.H. Diamond, J.Q. Dong, G.Y. Fu, A. Fukuyama, T.S. Hahm, D.E. Keyes, Y. Kishimoto, S. Klasky, L.L. Lao, K. Li, Z. Lin, B. Ludaescher, J. Manickam, N. Nakajima, T. Ozeki, N. Podhorszki, W.M. Tang, M.A. Vouk, R.E. Waltz, S.J. Wang, H. R. Wilson, X.Q. Xu, M. Yagi and F. Zonca. Simulation of Fusion Plasmas: Current Status and Future Direction. *Plasma Science and Technology*. 9:312-387, 2007.
81. Zhe Wang, Qin Lv, William Josephson, Wei Dong, Moses Charikar, Kai Li. Sizing Sketches: A Rank-Based Analysis for Similarity Search. *Proceedings of ACM SIGMETRICS 2007*, San Diego, CA, USA. June 2007.
82. Kai Li, Matthew Hibbs, Grant Wallace Maitreya Dunham, Rachel Sealfon, and Olga Troyanskaya. Scalable, Dynamic Analysis and Visualization for Genomic Datasets. *Proceedings of IPDPS Workshop on Next Generation Software*. March, 2007.
83. Lars A Bongo, Grant Wallace, Tore Larsen, Kai Li, Olga Troyanskaya. Systems Support for Remote Visualization of Genomics Applications over Wide Area Networks. *Proceedings of the International Workshop on Distributed, High-Performance and Grid Computing in Computational Biology (GCCB)*, 2006. Also in *Lecture Notes in Computer Science*, Springer Berlin - Heidelberg, vol 4360, pp157-174. March 7, 2007,

84. Qin Lv, William Josephson, Zhe Wang, Moses Charikar, and Kai Li, Efficient Filtering with Sketches in the Ferret Toolkit, In *Proceedings of the 8th ACM SIGMM International Workshop on Multimedia Information Retrieval*, October 2006.
85. Zhe Wang, Matthew Hoffman, Perry Cook and Kai Li. VFerret: Content-Based Similarity Search Tool for Continuous Archived Video. In *Proceedings of ACM workshop on Continuous Archiving and Recording of Personal Experiences CARPE-06*. October 2006.
86. Qin Lv, William Josephson, Zhe Wang, Moses Charikar, and Kai Li. Ferret: A Toolkit for Content-Based Similarity Search. In *Proceedings of ACM SIGOS EuroSys Conference*. April 2006.
87. Han Chen, Kai Li and Bin Wei. Memory Performance Optimizations For Real-Time Software HDTV Decoding. *Journal of VLSI Signal Processing*. 41(2): 193-207. September 2005.
88. Grant Wallace, Otto Anshus, Peng Bi, Han Chen, Yuqun Chen, Perry Cook, Adam Finkelstein, Thomas Funkhouser, Anoop Gupta, Matthew Hibbs, Kai Li, Zhiyan Liu, Rudrajit Samanta, Rahul Sukthankar, and Olga Troyanskaya. Tools and Applications for Large-Scale Display Walls. *IEEE Computer Graphics & Applications, Special Issue on Large Displays*. 25(4):24-33, July/August 2005.
89. Nitin Garg, Sumeet Sobti, Junwen Lai, Fengzhou Zheng, Kai Li, Arvind Krishnamurthy, and Randolph Wang. Bridging the Digital Divide: Storage Media + Postal Network = Generic High-Bandwidth Communication. *ACM Transactions on Storage (TOS)*. 1(2):246-275. May 2005.
90. Matthew Hibbs, Nathaniel C. Dirksen NC, Kai Li, and Olga G. Troyanskaya, Visualization Methods for Statistical Analysis of Microarray Clusters, *BMC Bioinformatics*, 6:115, May 2005.
91. Kai Li, Matthew Hibbs, Grant Wallace and Olga Troyanskaya. Dynamic Scalable Visualization for Collaborative Scientific Applications. *Proceedings of The Next Generation Software Workshop*. Denver, Colorado, April 2005.
92. Yuanyuan Zhou, Angelos Bilas, Suresh Jagannathan, Dimitrios Xinidis, Cezary Dubnicki and Kai Li. VI-attached Database Storage. *IEEE Transactions on Parallel and Distributed Systems*. 16(1): 35-50, January 2005.
93. Grant Wallace, Han Chen and Kai Li. Automatic Alignment of Tiled Displays for a Desktop Environment. *Journal of Software*. 15(12): 1776-1786. December 2004.
94. Qin Lv, Moses Charikar and Kai Li. Image Similarity Search with Compact Data Structures. *Proceedings of the ACM SIGIR 13th Conference on Information and Knowledge Management (CIKM)*. Pages 208-217, November 2004.
95. Wen Xu, Sanjeev Kumar, Kai Li. Fast Paths in Concurrent Programs. In *Proceedings of ACM/IEEE International Conference on Parallel Architecture and Compilation Techniques (PACT)*, Nice, France, September 2004.
96. Yuanyuan Zhou, Zhifeng Chen and Kai Li. Second-Level Buffer Cache Management. *IEEE Transactions on Parallel and Distributed Systems*. (15(6):505-519. June 2004.
97. Yuqun Chen, James Plank, and Kai Li. CLIP: A Checkpointing Tool for Message Passing Parallel Computers. *Scalable Input/Output: Achieving System Balance* (Edited by Daniel Reed). MIT Press. January 2004.
98. Yuanyuan Zhou, Limin Wang, Douglas W. Clark, and Kai Li. Thread Scheduling for Out-of-Core Applications with a Memory Server. *Scalable Input/Output: Achieving System Balance* (Edited by Daniel Reed). MIT Press. January 2004.
99. Yuanyuan Zhou, Liviu Iftode, and Kai Li. A Scalability Study of Shared Virtual Memory Systems. *Scalable Input/Output: Achieving System Balance* (Edited by Daniel Reed). MIT Press. January 2004.

100. Peter F. Corbett, Jean-Pierre Prost, Chris Demetriou, Garth Gibson, Erik Riedel, Jim Zelenka, Yuqun Chen, Ed Felten, Kai Li, John Hartman, Larry Peterson, Brian Bershad, Alec Wolman and Ruth Aydt. Proposal for a Common Parallel File System Programming Interface. *Scalable Input/Output: Achieving System Balance* (Edited by Daniel Reed). MIT Press. January 2004.
101. Grant Wallace, Han Chen, and Kai Li. DeskAlign: Automatically Aligning a Tiled Windows Desktop. *IEEE International Workshop on Projector-Camera Systems (PROCAMS)*, October 2003.
102. Zhifeng Chen, Yuanyuan Zhou and Kai Li. Eviction-based Placement for Storage Caches. In *Proceedings of USENIX Technical Conference*, June 2003.
103. Grant Wallace, Han Chen, and Kai Li.. Color Gamut Matching for Tiled Display Walls. *Immersive Projection Technology Workshop (IPT2003)*, May 2003.
104. Sanjeev Kumar, Kai Li. Using Model Checking to Debug Device Firmware. In *Proceedings of USENIX Operating Systems Design and Implementation (OSDI)*, 2002.
105. Han Chen, Rahul Sukthankar, Grant Wallace and Kai Li. Scalable Alignment of Large-Format Multi-Projector Displays Using Camera Homography Trees. In *Proceedings of IEEE Visualization (Vis2002)*, October 2002.
106. Han Chen, Kai Li, and Bin Wei. Memory Performance Optimizations for Real-Time Software HDTV Decoding. *IEEE International Conference on Multimedia and Expo (CME2002)*, August 2002.
107. Qin Lv, Pei Cao, Edith Cohen, Kai Li, and Scott Shenker, Search and Replication in Unstructured Peer-to-Peer Networks. *Proceedings of the ACM 16th International Conference on Supercomputing*, pages 84-95, June 2002.
108. Sanjeev Kumar, Kai Li. Dynamic Memory Management for Programmable Devices. *International Symposium of Memory Management (ISMM)*, pages 139-149. June 2002.
109. Yuanyuan Zhou, Suresh Jagannathan, Angelos Bilas, Cezary Dubnicki, James F Philbin, Kai Li. Experience with VI Communication for Database Storage. In *Proceedings of the 29th ACM International Symposium of Computer Architecture*. May 2002.
110. Han Chen, Kai Li, and Bin Wei. A Parallel Ultra-High Resolution MPEG-2 Video Decoder for PC Cluster Based Tiled Display System. In *Proceedings of International Parallel and Distributed Processing Symposium (IPDPS2002)*, April 2002.
111. Zhiyan Liu, Adam Finkelstein and Kai Li. Improving progressive view-dependent isosurface propagation. *Computers and Graphics*, 26 (2): 209-218. *Special Issue on Visualization on Very Large Datasets*. April 2002.
112. Han Chen, Grant Wallace, Anoop Gupta, and Kai Li, Tom Funkhouser, Perry Cook, Experiences with Scalability of Display Walls. In *Proceedings of the 7th Annual Immersive Projection Technology Symposium (IPT)*, March 2002.
113. Sanjeev Kumar, and Kai Li. Performance Impact of Using ESP to Implement VMMC Firmware. *Workshop on Novel Uses of System Area Networks (SAN-1)*, February 2002.
114. Rudrajit Samanta, Thomas Funkhouser, and Kai Li. Parallel Rendering with K-Way Replication. *IEEE Symposium on Parallel and Large-Data Visualization and Graphics*, October, 2001.
115. Han Chen, Yuqun Chen, Adam Finkelstein, Thomas Funkhouser, Kai Li, Zhiyan Liu, Rudrajit Samanta, and Grant Wallace. Data Distribution Strategies for High-Resolution Displays. *Computers & Graphics*, Special Issue on Mixed Realities - Beyond Conventions, 25(5):811-818. October 2001.
116. Sanjeev Kumar, Yitzhak Mandelbaum, Xiang Yu, Kai Li. ESP: A language for programmable devices. *Proceedings of ACM SIGPLAN Programming Language Design and Implementation (PLDI)*. June 2001.

117. Yuanyuan Zhou, James F. Philbin, and Kai Li. The Multi-Queue Replacement Algorithm for Second Level Buffer Caches. In *Proceedings of USENIX Technical Conference*, June 2001.
118. Yuqun Chen, Han Chen, Douglas W. Clark, Zhiyan Liu, Grant Wallace, and Kai Li. Software Environments for Cluster-based Display Systems (2001). The First IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid 2001), Brisbane, Australia, 15-18 May 2001.
119. Zhiyan Liu, Adam Finkelstein, and Kai Li. Progressive View-Dependent Isosurface Propagation. IEEE TCVG Symposium on Visualization (VisSym 2001). Ascona, Switzerland. May 28 - May 30, 2001.
120. Xiang Yu, Ben Gum, Yuqun Chen, Randy Wang, Kai Li, Arvind Krishnamurthy, and Thomas Anderson, Trading Capacity For Performance In A Disk Array. The 4th USENIX Symposium on Operating Systems Design and Implementation (OSDI 2000), Paradise Point Resort, San Diego, California , October 23-25, 2000.
121. Yuqun Chen, Douglas W. Clark, Adam Finkelstein, Timothy Housel, and Kai Li, Automatic Alignment Of High-Resolution Multi-Projector Displays Using An Uncalibrated Camera, IEEE Visualization 2000, Salt Lake City, Utah, October 8-13, 2000.
122. Rudrajit Samanta, Thomas Funkhouser, Kai Li, and Jaswinder Pal Singh, Hybrid Sort-First and Sort-Last Parallel Rendering with a Cluster of PCs. In *Proceedings of SIGGRAPH/Eurographics Workshop on Graphics Hardware*, August, 2000.
123. Kai Li, Han Chen, Yuqun Chen, Douglas W. Clark, Perry Cook, Stefanos Damianakis, Georg Essl, Adam Finkelstein, Thomas Funkhouser, Timothy Housel, Allison Klein, Zhiyan Liu, Emil Praun, Rudrajit Samanta, Ben Shedd, Jaswinder Pal Singh, George Tzanetakis, and Jiannan Zheng, Building and Using a Scalable Display Wall System, *IEEE Computer Graphics and Applications*, 20(4): 29-37, July 2000.
124. Thomas Funkhouser and Kai Li, Large Format Displays. *IEEE Computer Graphics and Applications*, 20(4): 20-21, July 2000.
125. Rudrajit Samanta, Thomas Funkhouser, Kai Li, and Jaswinder Pal Singh, Sort-First Parallel Rendering with a Cluster of PCs, *Technical Sketch in SIGGRAPH 2000*, July, 2000.
126. Rudro Samanta, Jiannan Zheng, Thomas Funkhouser, Kai Li, and Jaswinder Pal Singh. Load Balancing for Multi-Projector Rendering Systems, In *Proceedings of SIGGRAPH/Eurographics Workshop on Graphics Hardware*, August, 1999.
127. Yuqun Chen, Stefanos N. Damianakis, Sanjeev Kumar, Xiang Yu, and Kai Li. Porting a User-level Communication Architecture to NT: Experience and Performance (1999). In *Proceedings of 3rd Usenix Windows NT Symposium*, Seattle, Washington, July 12-15, 1999.
128. Liviu Iftode, Matthias Blumrich, Cezary Dubnicki, David L. Oppenheimer, Jaswinder Pal Singh, and Kai Li. Shared Virtual Memory with Automatic Update Support. In *Proceedings of the International Conference on Supercomputing*, Rhodes, Greece, June 1999.
129. Yuanyuan Zhou, Peter Chen, and Kai Li. Fast Cluster Failover Using Virtual Memory-Mapped Communication. *Proceedings of the International Conference on Supercomputing*, Rhodes, Greece, June 1999.
130. James S. Plank, Yuqun Chen, Kai Li, Micah Beck and Gerry Kingsley. Memory Exclusion: Optimizing the Performance of Checkpointing Systems (1999). *Software -- Practice and Experience*, Vol. 29, No. 2, pp. 125-142, 1999.
131. Y. Zhou, L. Wang and K. Li, Thread Scheduling for Out-of-Core Applications with Memory Server on Multicomputers. *Proceedings of the 6th Workshop on Input/Output in Parallel and Distributed Systems (IOPADS)*, May 1999.

132. Larry Peterson, Scott Karlin and Kai Li, OS Support for General Purpose Routers. ACM SIGOPS HotOS Workshop. March 1999.
133. James S. Plank, Kai Li and Michael A. Puening, Diskless Checkpointing. *IEEE Transactions on Parallel and Distributed Systems*, 9(10): 972-986, October, 1998.
134. Yuqun Chen, Czarek Dubnicki, Stefanos Damianakis, Angelos Bilas, and Kai Li. UTLB: A Mechanism for Translations on Network Interface. In *Proceedings of ACM Architectural Support for Programming Languages and Operating Systems (ASPLOS-VIII)*, San Jose, California, October 1998, pages 193-204.
135. Matthias A. Blumrich, Kai Li, Richard D. Alpert, Cezary Dubnicki, Edward W. Felten, Jonathan Sandberg. Retrospective: Virtual Memory-Mapped Network Interface for the SHRIMP Multicomputer. In *ACM 25 Years of the International Symposia on Computer Architecture, Selected Papers*. Edited by Gurindar Sohi, 1998. Pages 92-94.
136. Matthias Blumrich, Richard Alpert, Yuqun Chen, Douglas Clark, Stefanos Damianakis, Cezary Dubnicki, Edward Felten, Liviu Iftode, Kai Li, Margaret Martonosi, and Richard Shillner. Design Choices in the SHRIMP System: An Empirical Study. In *Proceedings of the ACM/IEEE 25th International Symposium on Computer Architecture*, Spain, June 1998, pages 330-341.
137. Bin Wei, Douglas Clark, Edward Felten, Kai Li, and Gordon Stoll. Performance Issues of a Distributed Frame Buffer on a Multicomputer. In *Proceedings of the 1998 ACM Eurographics/SIGGRAPH Workshop on Graphics Hardware*, Lisbon, Portugal, Pages 87-96, August 1998.
138. Cezary Dubnicki, Angelos Bilas, Yuqun Chen, Stefanos Damianakis and Kai Li. Shrimp Project Update: Myrinet Communication. *IEEE MICRO*, 18(1): 50-52. January 1998.
139. Minwen Ji, Edward Felten, and Kai Li. Performance Measurements for Multithreaded Programs. In *Proceedings of ACM SIGMETRICS / Performance '98: Joint International Conference on Measurement and Modeling of Computer Systems*. Pages 161-170. August 1998.
140. Cezary Dubnicki, Angelos Bilas, Yuqun Chen, Stefanos Damianakis and Kai Li. VMMC-2: Efficient Support for Reliable, Connection-Oriented Communication. *IEEE Hot Interconnects V*. August 1997.
141. Yuqun Chen, James S. Plank, and Kai Li. CLIP: A Checkpointing Tool for Message-Passing Parallel Programs. In *Proceedings of SC97: High Performance Networking & Computing*, San Jose, California, November 1997.
142. Yuanyuan Zhou, Liviu Iftode, Jaswinder Pal Singh, Kai Li, Brian R. Toonen, Ioannis Schoinas, Mark D. Hill, and David A. Wood. Relaxed Consistency and Coherence Granularity in DSM Systems: A Performance Evaluation. In *Proceedings of 6th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, Pages 193-205, June 1997.
143. Cezary Dubnicki, Angelos Bilas, Kai Li and James Philbin. Design and Implementation of Virtual Memory-Mapped Communication on Myrinet. In *Proceedings of the IEEE 11th International Parallel Processing Symposium*. Pages 388-396, April 1997.
144. Pei Cao, Edward W. Felten, Anna R. Karlin, and Kai Li. Implementation and Performance of Integrated Application-Controlled File Caching, Prefetching and Disk Scheduling. *ACM Transactions on Computer Systems*. 14(4):311-343. Nov 1996.
145. Fred Dougliis, Frans Kaashoek, Kai Li, Brian Marsh and Joseph Tauber. Storage Alternatives for Mobile Computers. *Mobile Computing*, Kluwer Academic Publishers, 1996.
146. Yuanyuan Zhou, Liviu Iftode and Kai Li. Performance Evaluation of Two Home-Based Lazy Release Consistency Protocols for Shared Virtual Memory Systems. In *Proceedings of the ACM SIGOPS/USENIX 2nd Symposium on Operating System Design and Implemenation (OSDI)*, Pages 75-88, October 1996.

147. Tracy Kimbrel, Andrew Tomkins, R. Hugo Patterson, Brian Bershad, Pei Cao, Edward W. Felten, Garth A. Gibson, Anna R. Karlin, and Kai Li. A Trace-driven Comparison of Algorithms for Parallel Prefetching and Caching. In *Proceedings of the ACM SIGOPS/USENIX 2nd Symposium on Operating System Design and Implementation (OSDI)*, Pages 19-34. October 1996.
148. James Philbin, Jan Edler, Otto J. Anshus, Craig Douglas, and Kai Li. Thread Scheduling for Cache Locality. In *Proceedings of the 7th ACM Conference on Architectural Support for Programming Languages and Operating Systems*, Cambridge, Massachusetts, October 1996. Pages 60-73.
149. Kai Li. Network of PCs as High-Performance Servers: Requirements for A Communication Mechanism. In *Proceedings of the IEEE 3rd International Conference on Massively Parallel Processing using Optical Interconnections* (Invited paper). Pages 225-230. October 1996.
150. Richard Alpert, Cezary Dubnicki, Edward Felten and Kai Li. Design and Implementation of NX Message Passing Using Shrimp Virtual Memory-Mapped Communication. In *Proceedings of the International Conference on Parallel Processing*. Pages 111-119. August 1996.
151. Liviu Iftode, Jaswinder Pal Singh and Kai Li. Scope Consistency: a Bridge between Release Consistency and Entry Consistency. In *Proceedings of the 8th Annual ACM Symposium on Parallel Algorithms and Architectures*. Pages 277-287. June 1996.
152. Liviu Iftode, Jaswinder Pal Singh and Kai Li. Understanding Application Performance on Shared Virtual Memory. In *Proceedings of the 23rd Annual International Symposium on Computer Architecture*. Pages 122-133. May 1996.
153. Edward Felten, Richard D. Alpert, Angelos Bilas, Matthias A. Blumrich, Douglas W. Clark, Stefanos Damianakis, Cezary Dubnicki, Liviu Iftode, and Kai Li. Early Experience with Message-Passing on the SHRIMP Multicomputer. In *Proceedings of the 23rd Annual International Symposium on Computer Architecture*. Pages 296-307. May 1996.
154. Cezary Dubnicki, Liviu Iftode, Edward W. Felten and Kai Li. Software Support for Virtual Memory-Mapped Communication. In *Proceedings of the IEEE 10th International Parallel Processing Symposium*. Pages 372-381. April 1996.
155. Matthias Blumrich, Cezary Dubnick, Edward Felten and Kai Li. Protected, User-Level DMA for the SHRIMP Network Interface. In *IEEE 2nd International Symposium on High-Performance Computer Architecture*. February 1996. Pages 154--165.
156. Liviu Iftode, Cezary Dubnick, Edward Felten and Kai Li. Improving Release-Consistent Shared Virtual Memory using Automatic Update. In *IEEE 2nd International Symposium on High-Performance Computer Architecture*. February 1996. Pages 14--25.
157. Bin Wei, Gordon Stoll, Douglas Clark, Edward Felten, Kai Li and Patrick Hanrahan. Synchronization for a Multi-Port Frame Buffer on a Mesh-Connected Multicomputer. In *IEEE Parallel Rendering Symposium '95*. Pages 81-88. October 1995.
158. Karin Petersen and Kai Li. Multiprocessor Cache Coherence Based on Virtual Memory Support. *Journal of Parallel and Distributed Computing*, 29(2):158-178, September 1995.
159. Pei Cao, Edward W. Felten, Anna Karlin and Kai Li. A study of Integrated Prefetching and Caching Strategies. In *Proceedings of the ACM SIGMETRICS*. Pages 188-197. 1995. **(Best paper award)**.
160. Gordon Stoll, Bin Wei, Douglas Clark, Edward Felten, Kai Li, and Patrick Hanrahan. Evaluating Multi-Port Frame Buffer Designs for a Mesh-Connected Multicomputer. In *Proceedings of the 22nd Annual International Symposium on Computer Architecture*. May 1995. Pages 96--105.

161. James S. Plank, Micah Beck, Gerry Kingsley, Kai Li. Libckpt: Transparent Checkpointing under Unix. In *Proceedings of the 1995 Winter USENIX Technical Conference*. New Orleans, LA, January 1995. Pages 213--223.
162. Matthias Blumrich, Cezary Dubnicki, Edward Felten, Kai Li and Malena Mesarina. Virtual-Memory-Mapped Network Interfaces. *IEEE MICRO*, 15(1):21--28. Feb 1995.
163. Pei Cao, Edward W. Felten and Kai Li. Implementation and Performance of Application-Controlled File Caching. In *Proceedings of the ACM SIGOPS/USENIX 1st Symposium on Operating Systems Design and Implementation (OSDI)*. Pages 165--178. November 1994.
164. Fred Douglass, Ramon Caceres, Frans Kaashoek, Kai Li, Brian Marsh, and Joshua A. Tauber. Storage Alternatives for Mobile Computers. In *Proceedings of the ACM SIGOPS/USENIX 1st Symposium on Operating Systems Design and Implementation (OSDI)*. November 1994. Pages 25--38.
165. James Plank and Kai Li. Performance Results of ICKP - A Consistent Checkpointer on the iPSC/860. *IEEE Parallel and Distributed Technologies*. 2(2):~62--67. Summer 1994.
166. Kai Li, Jeffrey Naughton and James Plank. Low-Latency Concurrent Checkpoint for Parallel Programs. *IEEE Transactions on Parallel and Distributed Computing*, 5(8):~874--879. 1994.
167. James S. Plank and Kai Li. Faster Checkpointing with $\{N+1\}$ Parity. In *Proceedings of IEEE 24th International Symposium on Fault-Tolerant Computing*. Pages 288--297. Austin, TX, June 1994.
168. Pei Cao, Edward Felten and Kai Li. Application-Controlled File Caching Policies. In *Proceedings of the 1994 Summer USENIX Technical Conference*. Pages 171--182. June 1994.
169. James S. Plank and Kai Li. Performance Results of $\{N+1\}$ Parity - A Consistent Checkpointer on the iPSC/860. In *Proceedings of Scalable High Performance Computing Conference*. Pages 686--693. Knoxville, TN, May, 1994.
170. Cezary Dubnicki, Kai Li and Malena Mesarina. Network Interface Support for User-Level Buffer Management. *Parallel Computer Routing and Communication. Lecture Notes in Computer Science (LCS 853)*, Edited by K. Bolding and L. Snyder. Springer-Verlag, April 1994.
171. Matthias Blumrich, Kai Li, Richard Alpert, Cezary Dubnicki, Edward Felten, and Jonathan Sandberg. Virtual Memory Mapped Network Interface for the Shrimp Multicomputer. In *ACM/IEEE Proceedings of the 21st Annual International Symposium on Computer Architecture*. Pages 142--153. April 1994. **Selected as one of the 43 most influential papers of 25 years of ISCA, 1998.**
172. Karin Petersen and Kai Li. An Evaluation of Multiprocessor Cache Coherence Based on Virtual Memory Support. In *Proceedings of the IEEE 8th International Parallel Processing Symposium*. pages 158-164. April 1994.
173. Ramon Caceres, Fred Douglass, Kai Li and Brian Marsh. Operating Systems Implications of Solid State Mobile Computers. In *Proceedings of the Fourth Workshop on Workstation Operating Systems (WWOS-IV)*. Pages 21--27. October 1993.
174. Karin Petersen and Kai Li. Cache Coherence for Shared Memory Multiprocessors Based on Virtual Memory Support. In *Proceedings of the IEEE 7th International Parallel Processing Symposium*. April 1993.
175. E. Barszcz, D. Black, D. Culler, H. Gordon, S. Groom, D. Kopetzky, R. Lee, K. Li, L. Lane, E. Lazowska, R. Light, J. Mukerji, E. Upchurch, and M. Wan. Operating Systems (chapter 7). *System Software and Tools for High Performance Computing Environments*. Edited by Paul Messina and Thomas Sterling, Society for Industrial and Applied Mathematics. Pages 75--92. 1993.

176. Songnian Zhou, Michael Stumm, Kai Li and David Wortman. Heterogeneous Distributed Shared Memory: An Experimental Study. *IEEE Transactions on Parallel and Distributed Computing*, 3(5):~540--554, 1992. (Also appeared in Multiprocessor Performance Measurement and Evaluation, edited by Laxmi N. Bhuyan and Xiaodong Zhang. IEEE Computer Society, 1994).
177. Liviu Iftode, Kai Li and Karin Petersen. Memory Servers for Multicomputers. In *Proceedings of the IEEE Spring CompCon'93*. Pages 539--547. February 1993.
178. Anne Rogers and Kai Li. Software Support for Speculative Load. In Proceedings of the ACM 5th International Conference on Architectural Support for Programming Languages and Operating Systems. Pages 38--50. 1992.
179. Kai Li. Scalability Issues of Shared Virtual Memory on Multicomputers. *Scalable Shared Memory Multiprocessors*, Edited by Michel Dubois and Shreekanth Thakkar, Kluwer Academic Publishers. 1992.
180. Kai Li, Jeffrey Naughton and James Plank. An Efficient Checkpointing Method for Multicomputers with Wormhole Routing. *International Journal of Parallel Programming*. 20(3):~159--180, 1991.
181. Anna Karlin, Kai Li, Mark Manasse, and Susan Owicki. Empirical Studies of Competitive Spinning on A Shared-Memory Multiprocessor. In *Proceedings of the 12th ACM Symposium on Operating Systems Principles*. October 1991.
182. Kai Li, Jeffrey Naughton and James Plank. Checkpointing Multicomputer Applications. In *Proceedings of the IEEE 10th Symposium on Reliable Distributed Systems*. Pisa, Italy. September 1991.
183. Kai Li and Karin Petersen. Evaluation of Extended Memory Systems. In Proceedings of the ACM/IEEE 18th Annual International Symposium on Computer Architecture. Pages 84--95. June 1991.
184. Andrew Appel and Kai Li. Virtual Memory Primitives for User Programs. In Proceedings of the ACM 4th International Conference on Architectural Support for Programming Languages and Operating Systems. pages 96--107, April 1991.
185. Kai Li, Jeffrey Naughton and James Plank. Real-time, Concurrent Checkpoint for Parallel Programs. In *Proceedings of the Second ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. Pages 79--88, 1990.
186. Kai Li and Richard Schaefer. Shared Virtual Memory for a Hypercube Multiprocessor. In *Proceedings of the 4th Conference on Hypercube Concurrent Computers and Applications*, March 1989.
187. Kai Li and Richard Schaefer. A Hypercube Shared Virtual Memory System. In *Proceedings of the 1989 International Parallel Processing Conference, Vol. I Architecture*, pages I-125--I132, August 1989.
188. Mark Greenstreet, Kai Li and J. Staunstrup. Synchronized Transitions on Multiprocessors. In *Proceedings of the 22rd Hawaii International Conference on System Sciences*. Pages 789--797, January 1989.
189. Kai Li and Paul Hudak. Memory Coherence in Shared Virtual Memory Systems. *ACM Transactions on Computer Systems*, 7(4):321--359, November 1989.
190. Kai Li and Jeffrey Naughton. Multiprocessor Main Memory Transaction Processing. In Proceedings of IEEE International Symposium on Databases in Parallel and Distributed Systems. Pages 177--187, December 1988.
191. Kai Li. IVY: A Prototype Shared Virtual Memory System for Parallel Computing. In *Proceedings of the 1988 International Conference on Parallel Computing. Vol. II Software*. Pages 94--101. August 1988.
192. Andrew Appel, John Ellis and Kai Li. Real-time Concurrent Collection on Stock Multiprocessors. In *Proceedings of ACM SIGPLAN '88 Conference on Programming*

Language Design and Implementation. Pages 11--20. June 1988. **Selected as one of the 50 most influential papers of 20 years of PLDI, 2004.**

193. Kai Li and Paul Hudak. Memory Coherence in Shared Virtual Memory Systems. In *Proceedings of the 5th ACM Conference on Principles of Distributed Computing*. Pages~229--239. August 1986.
194. Kai Li and Paul Hudak. A New List Compaction Method. *Software Practice and Experience*, 16(2):145--163, February 1986.