

David G. Andersen

Computer Science Department
Carnegie Mellon University
5000 Forbes Ave
Pittsburgh, PA 15213

Phone: (412) 268-3064
Fax: (412) 268-5576
dga+@cs.cmu.edu
<http://www.cs.cmu.edu/~dga/>

Education	MASSACHUSETTS INSTITUTE OF TECHNOLOGY	Cambridge, MA
	Ph.D. in Computer Science. February 2005. Thesis: “Improving End-to-End Availability using Overlay Networks” Minor: Computational Biology S.M. in Computer Science, 2001 <i>Advisor:</i> Hari Balakrishnan	
	UNIVERSITY OF UTAH	Salt Lake City, UT
	Bachelor of Science in Computer Science. <i>Cum Laude</i> , 1998 Bachelor of Science in Biology. <i>Cum Laude</i> , 1998	

Research Interests

Energy-efficient computing; distributed systems; computer networks.

Professional Experience

2019–	CTO and co-founder	Enriched Ag
2018–	Professor	Carnegie Mellon University Department of Computer Science A summary of my research activities at CMU and elsewhere begins on page 6.
2011–	Associate Professor	Carnegie Mellon University Department of Computer Science 2016 to 2019: Faculty co-director of the Intel Center for Science and Technology for Visual Cloud Computing.
2015–2019	Visiting Scientist	Google Brain
2005–2011	Assistant Professor	Carnegie Mellon University Department of Computer Science
1999–2004	Research Assistant	MIT Research assistant at the Laboratory for Computer Science (LCS / CSAIL). Worked in cooperation with the University of Utah on the RON+Emulab testbed. Major projects at MIT include Resilient Overlay Networks (RON), Multihomed Overlay Networks (MONET), Mayday, and the Congestion Manager.
Summer 2001	Intern	Compaq SRC Summer internship working on the Secure Network Attached Disks project.
1997-1999	Research Assistant / Research Associate	University of Utah One year as an undergraduate and one year as a staff research associate in the Flux research group at the University of Utah.
1996-1997	Research Assistant	Department of Biology, University of Utah Undergraduate research assistantship in the Wayne Potts Laboratory in the Department of Biology.
1995-1997	Co-founder and CTO, ArosNet, Inc.	Acted in a directorial and technical capacity over technical operations: network design and topology planning, software development, consulting projects, and short-term research. During my three years with the company, ArosNet grew from its inception to become the third largest ISP in Utah.

Consultant and Expert Witness *Multi-State Lottery Association, Intel, Banner & Witcoff, LLC, IJNT, Inc., Ascensus, others.*

Provided network design, security, and intellectual property consulting and expert witnessing. Research consulting for Intel Research.

Refereed Publications

- [1] Conglong Li, Matthew K. Mukerjee, David G. Andersen, Srinivasan Seshan, Michael Kaminsky, George Porter, and Alex C. Snoeren. Using indirect routing to recover from network traffic scheduling estimation error. In *ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS)*, May 2017.
- [2] Anuj Kalia, Michael Kaminsky, and David G. Andersen. FaSST: Fast, scalable and simple distributed transactions with two-sided RDMA datagram RPCs. In *Proc. 12th USENIX OSDI*, Savannah, GA, November 2016.
- [3] Anuj Kalia, Michael Kaminsky, and David G. Andersen. Design guidelines for high-performance RDMA systems. In *Proc. USENIX Annual Technical Conference*, Denver, CO, June 2016.
- [4] Huanchen Zhang, David G. Andersen, Andrew Pavlo, Michael Kaminsky, Lin Ma, and Rui Shen. Reducing the storage overhead of main-memory OLTP databases with hybrid indexes. In *Proc. ACM SIGMOD*, San Francisco, USA, June 2016.
- [5] Sheng Li, Hyeontaek Lim, Victor W. Lee, Jung Ho Ahn, Anuj Kalia, Michael Kaminsky, David G. Andersen, O. Seongil, Sukhan Lee, and Pradeep Dubey. Achieving a billion key-value requests per second on a single server. *IEEE Micro's top picks from the computer architecture conferences 2016*, 36(3), may/june 2016.
- [6] Sheng Li, Hyeontaek Lim, Victor W. Lee, Jung Ho Ahn, Anuj Kalia, Michael Kaminsky, David G. Andersen, O. Seongil, Sukhan Lee, and Pradeep Dubey. Full-stack architecting to achieve a billion requests per second throughput on a single key-value store server platform. *ACM Transactions on Computer Systems (TOCS)*, 34(2), April 2016.
- [7] Xiaozhou Li, Raghav Sethi, Michael Kaminsky, David G. Andersen, and Michael J. Freedman. Be fast, cheap and in control with switchkv. In *Proc. 13th USENIX NSDI*, Santa Clara, CA, May 2016.
- [8] Hyeontaek Lim, David G. Andersen, and Michael Kaminsky. Towards accurate and fast evaluation of multi-stage log-structured designs. In *Proc. USENIX Conference on File and Storage Technologies*, Santa Clara, CA, February 2016.
- [9] He Liu, Matthew K. Mukerjee, Conglong Li, Nicolas Feltman, George Papen, Stefan Savage, Srinivasan Seshan, Geoffrey M. Voelker, David G. Andersen, Michael Kaminsky, George Porter, and Alex C. Snoeren. Scheduling techniques for hybrid circuit/packet networks. In *Proc. CoNEXT*, December 2015.
- [10] Dong Zhou, Bin Fan, Hyeontaek Lim, David G. Andersen, Michael Kaminsky, Michael D. Mitzenmacher, Ren Wang, and Ajaypal Singh. Scaling up clustered network appliances with ScaleBricks. In *Proc. ACM SIGCOMM*, London, UK, August 2015.
- [11] Sheng Li, Hyeontaek Lim, Victor W. Lee, Jung Ho Ahn, Anuj Kalia, Michael Kaminsky, David G. Andersen, O. Seongil, Sukhan Lee, and Pradeep Dubey. Architecting to achieve a billion requests per second throughput on a single key-value store server platform. In *ISCA*, 2015.
- [12] Anuj Kalia, Dong Zhou, Michael Kaminsky, and David G. Andersen. Raising the bar for using GPUs in software packet processing. In *Proc. 12th USENIX NSDI*, Oakland, CA, May 2015.
- [13] Iulian Moraru, David G. Andersen, and Michael Kaminsky. Paxos quorum leases: Fast reads without sacrificing writes. In *Proc. 5th ACM Symposium on Cloud Computing (SOCC)*, Seattle, WA, November 2014.
- [14] Wyatt Lloyd, Michael J. Freedman, Michael Kaminsky, and David G. Andersen. Don't settle for eventual consistency. *Communications of the ACM*, 57(5), May 2014.

- [15] Bin Fan, David G. Andersen, Michael Kaminsky, and Michael D. Mitzenmacher. Cuckoo filter: Practically better than Bloom. In *Proc. CoNEXT*, December 2014.
- [16] Mu Li, David G. Andersen, Jun Woo Park, Alexander J. Smola, Amr Ahmed, Vanja Josifovski, James Long, Eugene J. Shekita, and Bor-Yiing Su. Scaling distributed machine learning with the parameter server. In *Proc. 11th USENIX OSDI*, Broomfield, CO, October 2014.
- [17] Anuj Kalia, Michael Kaminsky, and David G. Andersen. Using RDMA efficiently for key-value services. In *Proc. ACM SIGCOMM*, Chicago, IL, August 2014.
- [18] David Naylor, Matthew K. Mukerjee, et al. XIA: Architecting a more trustworthy and evolvable Internet. *ACM Computer Communications Review*, 44(3), July 2014.
- [19] Hyeontaek Lim, Dongsu Han, David G. Andersen, and Michael Kaminsky. MICA: A holistic approach to fast in-memory key-value storage. In *Proc. 11th USENIX NSDI*, Seattle, WA, April 2014.
- [20] Xiaozhou Li, David G. Andersen, Michael Kaminsky, and Michael J. Freedman. Algorithmic improvements for fast concurrent cuckoo hashing. In *Proc. 9th ACM European Conference on Computer Systems (EuroSys)*, April 2014.
- [21] Hrishikesh Amur, Wolfgang Richter, David G. Andersen, Michael Kaminsky, Karsten Schwan, Athula Balanachandran, and Erik Zawadzki. Memory-efficient GroupBy-Aggregate using compressed buffer trees. In *Proc. 4rd ACM Symposium on Cloud Computing (SOCC)*, Santa Clara, CA, October 2013.
- [22] Dong Zhou, Bin Fan, Hyeontaek Lim, David G. Andersen, and Michael Kaminsky. Scalable, High Performance Ethernet Forwarding with CuckooSwitch. In *Proc. 9th International Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, December 2013.
- [23] Iulian Moraru, David G. Andersen, Michael Kaminsky, Parthasarathy Ranganathan, Niraj Tolia, and Nathan Binkert. Consistent, durable, and safe memory management for byte-addressable non volatile main memory. In *Proceedings of first Conference on Timely Results in Operating Systems TRIOS*, November 2013.
- [24] Iulian Moraru, David G. Andersen, and Michael Kaminsky. There is more consensus in egalitarian parliaments. In *Proc. 24th ACM Symposium on Operating Systems Principles (SOSP)*, Farmington, PA, November 2013.
- [25] Bin Fan, Dong Zhou, Hyeontaek Lim, Michael Kaminsky, and David G. Andersen. When cycles are cheap, some tables can be huge. In *Proc. HotOS XIV*, Santa Ana Pueblo, NM, May 2013.
- [26] Dong Zhou, David G. Andersen, and Michael Kaminsky. Space-efficient, high-performance rank & select structures on uncompressed bit sequences. In *Proc. 12th International Symposium on Experimental Algorithms (SEA)*, June 2013.
- [27] Bin Fan, David G. Andersen, and Michael Kaminsky. MemC3: Compact and concurrent memcache with dumber caching and smarter hashing. In *Proc. 10th USENIX NSDI*, Lombard, IL, April 2013.
- [28] Wyatt Lloyd, Michael J. Freedman, Michael Kaminsky, and David G. Andersen. Stronger semantics for low-latency geo-replicated storage. In *Proc. 10th USENIX NSDI*, Lombard, IL, April 2013.
- [29] Hyeontaek Lim, David G. Andersen, and Michael Kaminsky. Practical batch-updatable external hashing with sorting. In *Proc. Meeting on Algorithm Engineering and Experiments (ALENEX)*, January 2013.
- [30] Amar Phanishayee, David G. Andersen, Himabindu Pucha, Anna Povzner, and Wendy Belluomini. Flex-KV: Enabling high-performance and flexible KV systems. In *Proc. Workshop on Management of Big Data Systems*, September 2012.
- [31] Vijay Vasudevan, Michael Kaminsky, and David G. Andersen. Using vector interfaces to deliver millions of IOPS from a networked key-value storage server. In *Proc. 3rd ACM Symposium on Cloud Computing (SOCC)*, San Jose, CA, October 2012.

- [32] Dongsu Han, Ashok Anand, Fahad Dogar, Boyan Li, Hyeontaek Lim, Michel Machado, Arvind Mukundan, Wenfei Wu, Aditya Akella, David G. Andersen, John W. Byers, Srinivasan Seshan, and Peter Steenkiste. XIA: Efficient support for evolvable internetworking. In *Proc. 9th USENIX NSDI*, San Jose, CA, April 2012.
- [33] Iulian Moraru and David G. Andersen. Exact pattern matching with feed-forward bloom filters. *Journal of Experimental Algorithmics (JEA)*, 17(1), July 2012.
- [34] Hyeontaek Lim, Bin Fan, David G. Andersen, and Michael Kaminsky. SILT: A memory-efficient, high-performance key-value store. In *Proc. 23rd ACM Symposium on Operating Systems Principles (SOSP)*, Cascais, Portugal, October 2011.
- [35] Wyatt Lloyd, Michael J. Freedman, Michael Kaminsky, and David G. Andersen. Don't settle for eventual: Scalable causal consistency for wide-area storage with COPS. In *Proc. 23rd ACM Symposium on Operating Systems Principles (SOSP)*, Cascais, Portugal, October 2011.
- [36] Bin Fan, Hyeontaek Lim, David G. Andersen, and Michael Kaminsky. Small cache, big effect: Provable load balancing for randomly partitioned cluster services. In *Proc. 2nd ACM Symposium on Cloud Computing (SOCC)*, Cascais, Portugal, October 2011.
- [37] Hamid Hajabdolali Bazzaz, Malveeka Tewari, Guohui Wang, George Porter, T. S. Eugene Ng, David G. Andersen, Michael Kaminsky, Michael A. Kozuch, and Amin Vahdat. Switching the optical divide: Fundamental challenges for hybrid electrical/optical datacenter networks. In *Proc. 2nd ACM Symposium on Cloud Computing (SOCC)*, Cascais, Portugal, October 2011.
- [38] Ashok Anand, Fahad R. Dogar, Dongsu Han, Boyan Li, Hyeontaek Lim, Michel Machado, Wenfei Wu, Aditya Akella, David G. Andersen, John W. Byers, Srinivasan Seshan, and Peter Steenkiste. XIA: An architecture for an evolvable and trustworthy Internet. In *Proc. ACM Hotnets-X*, Cambridge, MA, USA., November 2011.
- [39] Xin Zhang, Hsu-Chun Hsiao, Geoffrey Hasker, Haowen Chan, Adrian Perrig, and David G. Andersen. SCION: Scalability, control, and isolation on next-generation networks. In *Proc. IEEE Symposium on Security and Privacy*, Oakland, CA, May 2011.
- [40] David G. Andersen, Jason Franklin, Michael Kaminsky, Amar Phanishayee, Lawrence Tan, and Vijay Vasudevan. FAWN: A fast array of wimpy nodes. *Communications of the ACM*, 54(7):101–109, July 2011.
- [41] Vijay Vasudevan, David G. Andersen, and Michael Kaminsky. The case for VOS: The vector operating system. In *Proc. HotOS XIII*, Napa, CA, May 2011.
- [42] Anirudh Badam, Dongsu Han, David G. Andersen, Michael Kaminsky, Konstantina Papagiannaki, and Srinivasan Seshan. The hare and the tortoise: Taming wireless losses by exploiting wired reliability. In *Proc. ACM MobiHoc*, Paris, France, May 2011.
- [43] Vijay Vasudevan, David G. Andersen, Michael Kaminsky, Jason Franklin, Michael A. Kozuch, Iulian Moraru, Padmanabhan Pillali, and Lawrence Tan. Challenges and opportunities for efficient computing with FAWN. *Operating Systems Review*, (1):33–44, January 2011.
- [44] Iulian Moraru and David G. Andersen. Exact pattern matching with feed-forward bloom filters. In *Proceedings of the Workshop on Algorithm Engineering and Experiments (ALENEX11)*, ALENEX 2011. Society for Industrial and Applied Mathematics, January 2011.
- [45] Bin Fan, David G. Andersen, Michael Kaminsky, and Konstantina Papagiannaki. Balancing throughput, robustness, and in-order delivery in P2P VoD. In *Proc. CoNEXT*, December 2010.
- [46] Guohui Wang, David G. Andersen, Michael Kaminsky, Michael Kozuch, T. S. Eugene Ng, Konstantina Papagiannaki, and Michael Ryan. c-Through: Part-time optics in data centers. In *Proc. ACM SIGCOMM*, New Delhi, India, August 2010.
- [47] Rick McGeer, David G. Andersen, and Stephen Schwab. The network testbed mapping problem. In *Proc. 6th International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom)*, May 2010.

- [48] Sang Kil Cha, Iulian Moraru, Jiyong Jang, John Truelove, David Brumley, and David G. Andersen. SplitScreen: Enabling efficient, distributed malware detection. In *Proc. 7th USENIX NSDI*, San Jose, CA, April 2010.
- [49] Kanat Tangwongsan, Himabindu Pucha, David G. Andersen, and Michael Kaminsky. Efficient similarity estimation for systems exploiting data redundancy. In *Proc. IEEE INFOCOM*, San Diego, CA, March 2010.
- [50] David Sontag, Yang Zhang, Amar Phanishayee, David G. Andersen, and David Karger. Scaling all-pairs overlay routing. In *Proc. CoNEXT*, December 2009.
- [51] Guohui Wang, David G. Andersen, Michael Kaminsky, Michael Kozuch, T. S. Eugene Ng, Konstantina Papagiannaki, Madeleine Glick, and Lily Mummert. Your data center is a router: The case for reconfigurable optical circuit switched paths. In *Proc. ACM Hotnets-VIII*, New York City, NY, USA., October 2009.
- [52] David G. Andersen, Jason Franklin, Michael Kaminsky, Amar Phanishayee, Lawrence Tan, and Vijay Vasudevan. FAWN: A fast array of wimpy nodes. In *Proc. 22nd ACM Symposium on Operating Systems Principles (SOSP)*, Big Sky, MT, October 2009.
- [53] Vijay Vasudevan, Amar Phanishayee, Hiral Shah, Elie Krevat, David G. Andersen, Gregory R. Ganger, Garth A. Gibson, and Brian Mueller. Safe and effective fine-grained TCP retransmissions for datacenter communication. In *Proc. ACM SIGCOMM*, Barcelona, Spain, August 2009.
- [54] B. Aditya Prakash, Nicholas Valler, David G. Andersen, Michalis Faloutsos, and Christos Faloutsos. BGP-lens: patterns and anomalies in Internet routing updates. In *Proc. 15th SIGKDD International Conference On Knowledge Discovery and Data Mining, industrial track*, Paris, France, June 2009.
- [55] Bryan Parno, Jonathan M. McCune, Dan Wendlandt, David G. Andersen, and Adrian Perrig. CLAMP: Practical prevention of large-scale data leaks. In *Proc. IEEE Symposium on Security and Privacy*, Oakland, CA, May 2009.
- [56] Vijay Vasudevan, Jason Franklin, David Andersen, Amar Phanishayee, Lawrence Tan, Michael Kaminsky, and Iulian Moraru. FAWNdamementally power-efficient clusters. In *Proc. HotOS XII*, Monte Verita, Switzerland, May 2009.
- [57] Dongsu Han, David G. Andersen, Michael Kaminsky, Konstantina Papagiannaki, and Srinivasan Seshan. Access point localization using local signal strength gradient. In *Passive & Active Measurement (PAM)*, Seoul, South Korea, April 2009.
- [58] George Nychis, Vyas Sekar, David G. Andersen, Hyong Kim, and Hui Zhang. An empirical evaluation of entropy-based traffic anomaly detection. In *Proc. Internet Measurement Conference*, Vouliagmeni, Greece, October 2008.
- [59] Dongsu Han, Aditya Agarwala, David G. Andersen, Michael Kaminsky, Konstantina Papagiannaki, and Srinivasan Seshan. Mark-and-Sweep: Getting the “inside” scoop on neighborhood networks. In *Proc. Internet Measurement Conference*, Vouliagmeni, Greece, October 2008.
- [60] Fahad Dogar, Amar Phanishayee, Himabindu Pucha, Olatunji Ruwase, and David Andersen. Ditto - a system for opportunistic caching in multi-hop wireless mesh networks. In *Proc. ACM MobiCom*, San Francisco, CA, September 2008.
- [61] David G. Andersen, Hari Balakrishnan, Nick Feamster, Teemu Koponen, Daekyeong Moon, and Scott Shenker. Accountable Internet Protocol (AIP). In *Proc. ACM SIGCOMM*, Seattle, WA, August 2008.
- [62] Dan Wendlandt, David Andersen, and Adrian Perrig. Perspectives: Improving SSH-style host authentication with multi-path probing. In *Proc. USENIX Annual Technical Conference*, Boston, MA, June 2008.
- [63] Himabindu Pucha, Michael Kaminsky, David G. Andersen, and Michael A. Kozuch. Adaptive file transfers for diverse environments. In *Proc. USENIX Annual Technical Conference*, Boston, MA, June 2008.

- [64] Vyas Sekar, Michael K. Reiter, Walter Willinger, Hui Zhang, Ramana Rao Kompella, and David G. Andersen. cSamp: A system for network-wide flow monitoring. In *Proc. 5th USENIX NSDI*, San Francisco, CA, April 2008.
- [65] Mikhail Afanasyev, David G. Andersen, and Alex C. Snoeren. Efficiency through eavesdropping: Link-layer packet caching. In *Proc. 5th USENIX NSDI*, San Francisco, CA, April 2008.
- [66] Bryan Parno, Adrian Perrig, and David G. Andersen. SNAPP: Stateless network-authenticated path pinning. In *Proc. ACM Symposium on Information, Computer, and Communications Security (ASIACCS)*, Tokyo, Japan, March 2008.
- [67] Amar Phanishayee, Elie Krevat, Vijay Vasudevan, David G. Andersen, Gregory R. Ganger, Garth A. Gibson, and Srinivasan Seshan. Measurement and analysis of TCP throughput collapse in cluster-based storage systems. In *Proc. USENIX Conference on File and Storage Technologies (FAST)*, San Jose, CA, February 2008.
- [68] Elie Krevat, Vijay Vasudevan, Amar Phanishayee, David G. Andersen, Gregory R. Ganger, Garth A. Gibson, and Srinivasan Seshan. On application-level approaches to avoiding TCP throughput collapse in cluster-based storage systems. In *Proc. Petascale Data Storage Workshop at Supercomputing'07*, November 2007.
- [69] David Andersen, Hari Balakrishnan, Nick Feamster, Teemu Koponen, Daekyeong Moon, and Scott Shenker. Holding the Internet accountable. In *Proc. 6th ACM Workshop on Hot Topics in Networks (Hotnets-VI)*, Atlanta, GA, November 2007.
- [70] Matthew W. Dunlop, Ginger Perng, and David G. Andersen. SWAP: Shared wireless access protocol (using reciprocity). In *IEEE Workshop on Information Assurance*, June 2007.
- [71] Himabindu Pucha, David G. Andersen, and Michael Kaminsky. Exploiting similarity for multi-source downloads using file handprints. In *Proc. 4th USENIX NSDI*, Cambridge, MA, April 2007.
- [72] Dan Wendlandt, Ioannis Avramopoulos, David Andersen, and Jennifer Rexford. Don't secure routing protocols, secure data delivery. In *Proc. 5th ACM Workshop on Hot Topics in Networks (Hotnets-V)*, Irvine, CA, November 2006.
- [73] Niraj Tolia, Michael Kaminsky, David G. Andersen, and Swapnil Patil. An architecture for Internet data transfer. In *Proc. 3rd Symposium on Networked Systems Design and Implementation (NSDI)*, San Jose, CA, May 2006.
- [74] David G. Andersen, Hari Balakrishnan, M. Frans Kaashoek, and Rohit Rao. Improving Web availability for clients with MONET. In *Proc. 2nd USENIX NSDI*, Boston, MA, May 2005.
- [75] David G. Andersen, Alex C. Snoeren, and Hari Balakrishnan. Best-path vs. multi-path overlay routing. In *Proc. ACM SIGCOMM Internet Measurement Conference*, Miami, FL, October 2003.
- [76] Nick Feamster, David Andersen, Hari Balakrishnan, and M. Frans Kaashoek. Measuring the effects of Internet path faults on reactive routing. In *Proc. ACM SIGMETRICS*, San Diego, CA, June 2003.
- [77] Marcos K. Aguilera, Minwen Ji, Mark Lillibridge, John MacCormick, Erwin Oertli, David G. Andersen, Mike Burrows, Timothy Mann, and Chandramohan Thekkath. Block-Level Security for Network-Attached Disks. In *Proc. 2nd USENIX Conference on File and Storage Technologies*, March 2003.
- [78] David G. Andersen. Mayday: Distributed Filtering for Internet Services. In *Proc. 4th USENIX Symposium on Internet Technologies and Systems (USITS)*, Seattle, Washington, March 2003. PDF/ps updated 2008 to correct an unclear explanation.
- [79] David G. Andersen, Nick Feamster, Steve Bauer, and Hari Balakrishnan. Topology inference from BGP routing dynamics. In *Proc. ACM SIGCOMM Internet Measurement Workshop*, Marseille, France, November 2002.
- [80] David G. Andersen, Hari Balakrishnan, M. Frans Kaashoek, and Robert Morris. Resilient Overlay Networks. In *Proc. 18th ACM Symposium on Operating Systems Principles (SOSP)*, Banff, Canada, October 2001.

- [81] David G. Andersen, Hari Balakrishnan, M. Frans Kaashoek, and Robert Morris. The Case for Resilient Overlay Networks. In *Proc. HotOS VIII*, Schloss-Elmau, Germany, May 2001.
- [82] Alex Snoeren, David Andersen, and Hari Balakrishnan. Fine-Grained Failover Using Connection Migration. In *Proc. 3rd USENIX Symposium on Internet Technologies and Systems (USITS)*, San Francisco, CA, March 2001.
- [83] David Andersen, Deepak Bansal, Dorothy Curtis, Srinivasan Seshan, and Hari Balakrishnan. System support for bandwidth management and content adaptation in Internet applications. In *Proc. 4th USENIX OSDI*, pages 213 – 225, San Diego, CA, November 2000.
- [84] Ray Spencer, Stephen Smalley, Peter Loscocco, Mike Hibler, David Andersen, and Jay Lepreau. The Flask Security Architecture: System Support for Diverse Security Policies. In *Proc. 8th USENIX Security Symposium*, Washington, DC, August 1999.

Other Articles

- [85] Martin Abadi and David G. Andersen. Learning to protect communications with adversarial neural cryptography. *CoRR*, 2016.
- [86] Hrishikesh Amur, David G. Andersen, Michael Kaminsky, and Karsten Schwan. Design of a write-optimized data store. Technical Report GIT-CERCS-13-08, Georgia Tech CERCS, 2013.
- [87] Dongsu Han, David G. Andersen, Michael Kaminsky, Konstantina Papagiannaki, and Srinivasan Seshan. Hulu in the neighborhood. In *Proc. 3rd International Conference on Communication Systems and Networks (COMSNET)*, January 2011. invited paper.
- [88] David G. Andersen and Steven Swanson. Rethinking flash in the data center. *IEEE Micro*, 2010. (Invited commentary article).
- [89] Vijay Vasudevan, David G. Andersen, Michael Kaminsky, Lawrence Tan, Jason Franklin, and Iulian Moraru. Energy-efficient cluster computing with FAWN: Workloads and implications. In *Proc. e-Energy 2010*, Passau, Germany, April 2010. (invited paper).
- [90] Vijay Vasudevan, David G. Andersen, and Hui Zhang. On Internet availability: Where does path choice matter? Technical Report CMU-CS-TR-114, Carnegie Mellon University, March 2009.
- [91] Vijay Vasudevan, David G. Andersen, and Hui Zhang. Understanding the AS-level path disjointness provided by multi-homing. Technical Report CMU-CS-TR-141, Carnegie Mellon University, July 2007.
- [92] Iulian Moraru and David G. Andersen. Fast cache for your text: Accelerating exact pattern matching with feed-forward bloom filters. Technical Report CMU-CS-09-159, Department of Computer Science, Carnegie Mellon University, November 2009.
- [93] Madeleine Glick, David G. Andersen, Michael Kaminsky, and Lily Mummert. Dynamically reconfigurable optical links for high-bandwidth data center networks. In *Optical Fiber Comm. Conference (OFC)*, March 2009. (invited paper).
- [94] Szymon Jakubczak, David G. Andersen, Michael Kaminsky, Konstantina Papagiannaki, and Srinivasan Seshan. Link-alike: using wireless to share network resources in a neighborhood. *ACM SIGMOBILE MC2R*, 12(4), October 2008. (invited paper).
- [95] Elaine Shi, Ion Stoica, David Andersen, and Adrian Perrig. OverDoSe: A generic DDoS protection service using an overlay network. Technical Report CMU-CS-06-114, Carnegie Mellon University Computer Science Department, February 2006.
- [96] Niraj Tolia, David G. Andersen, and M. Satyanarayanan. Quantifying interactive user experience on thin clients. *IEEE Computer*, 39(3), March 2006.
- [97] David G. Andersen and Nick Feamster. Challenges and opportunities in Internet data mining. Technical Report CMU-PDL-06-102, Carnegie Mellon University, January 2006.
- [98] David G. Andersen. Critical networking infrastructure in a suitcase. In *NSF Workshop on Research Challenges in Distributed Computer Systems*, September 2005. (position paper).

- [99] David G. Andersen. Overlay networks: Networking on top of the network. ACM Computing Reviews Hot Topics essay - http://www.reviews.com/hottopic/hottopic_essay.cfm, September 2004.
- [100] David G. Andersen, Hari Balakrishnan, M. Frans Kaashoek, and Robert Morris. Experience with an Evolving Overlay Network Testbed. *ACM Computer Communications Review*, 33(3):13–19, July 2003.

All papers are available online at: <http://www.cs.cmu.edu/~dga/papers/>

Patents

“Method and system for managing access control” (number 20040243827, published) and “Method and system for securing block-based storage with capability data.” Marcos K. Aguilera, Minwen Ji, Mark Lillibridge, John MacCormick, Oerwin Oertli, Dave Andersen, Mike Burrows, Tim Mann, Chandu Thekkath. filed May 2003, (number 20040243828, Abandoned)

Software Artifacts

Memory efficient practical data structures - the libCuckoo open source release is the fastest available concurrent hash table for many workloads. The Cuckoo Filter outperforms Bloom filterse for many workloads.

Memory efficient data management - memory-efficient multi-pattern search; the MemC3 memory caching layer.

FAWN-KV A high-performance, log-structured key-value storage system designed to operate on Flash storage.

SplitScreen High-performance, open source virus scanning. Distributed as patches to the popular “ClamAV” virus scanner.

Perspectives The Firefox plugin and SSH patches for automatically authenticating self-signed certificates. This software has been installed by over 30,000 users.

DOT The Data-Oriented Transfer service: End-host software that provides a flexible, modular data transfer service on behalf of other applications.

CM The Congestion Manager: Congestion control software for end-hosts.

RON Resilient Overlay Networks: End-host based overlay routing that routes around failures and poor performance.

MONET A Web proxy derived from the Squid proxy that uses multiple local network providers and an overlay network of peer proxies to provide highly available and fast Web access.

The RON Testbed A 36-site Internet testbed used by a dozen or so external researchers, in addition to several researchers within MIT and Carnegie Mellon.

Talks

2012 “FAWN: From Architecture to Algorithms.” University of Wisconsin, Madison; University of California, San Diego; Harvard University.

2010 Intel Research, Beijing, China, “Experiences with FAWN.”

2010 Microsoft Research Asia, Beijing, China, “Experiences with FAWN.”

2006–2010 Yearly Cylab seminar and review meeting speaker

2010 Industry faculty visit days: Network Appliance, Microsoft, Google.

Apr 2010 Invited talk, IBM Almaden, “Experiences with FAWN.”

Apr 2010 UCSD workshop on non-volatile memory, “System and Algorithmic Adaptation for Flash”

Jun 2009 University of Washington mini-workshop on systems, “The Case for a Fast Array of Wimpy Nodes”

2009 Intel Open House, featured speaker

Jul 2009 Network Appliance, “A Fast Array of Wimpy Nodes”

May 2009	Massachusetts Institute of Technology, “The Case for a Fast Array of Wimpy Nodes”
Dec 2008	NIPS (Neural Information Processing Symposium) workshop on Parallel Implementations of Learning Algorithms, “Using a Fast Array of Wimpy Nodes.”
Aug 2008	SIGCOMM talk, “Accountable Internet Protocol.”
Nov 2006	Georgia Institute of Technology, “Building Better Data Transfers”
Nov 2006	University of Washington Database Seminar, “Easing the Pain of Network Data Analysis”
May 2006	Brown University, “Data-Oriented Transfer”
Sep 2005	MIT, NSF Grand Challenges in Networking Workshop
Jan 2005	NSF Workshop in Overcoming Barriers to Disruptive Innovation in Networking
May 2005	NSDI Conference paper talk, “Improving Web Availability for Clients with MONET”
Oct 2003	Internet Measurement Conference Conference paper talk, “Best-Path vs. Multi-Path Overlay Routing.”
Aug 2003	IRIS Student Workshop “An End-to-End Approach to Network Failure Recovery”
Mar 2003	USITS Conference paper talk, “Mayday: Distributed Filtering for Internet Services.”
Nov 2002	Internet Measurement Workshop Workshop paper talk, “Topology Inference from BGP Routing Dynamics.”
Oct 2002	OPENSIG invited talk, “RON: Choosing Resiliency.”
Oct 2001	SOSP Conference paper talk, “Resilient Overlay Networks.”
May 2001	HotOS VIII Workshop paper talk, “The Case for Resilient Overlay Networks.”
Oct 2000	OSDI Conference paper talk, “System Support for Bandwidth Management and Content Adaptation in Internet Applications.”
Feb 1999	OSDI works-in-progress “Agile: A hierarchically Extensible Security Policy Architecture.”

Selected Honors and Awards

2022	ACM SIGOPS Mark Weiser Award.
2020	Best paper award, ACM Symposium on Cloud Computing (SOCC’20).
2019	Best paper award, USENIX Symposium on Networked Systems Design and Implementation (NSDI).
2018	Best paper award, ACM SIGMOD International Conference on Management of Data (SIGMOD’18).
2016	Best paper award at USENIX ATC’16, IEEE Micro Top Picks 2016 award.
2014	Best paper award at ACM Symposium on Cloud Computing (SoCC).
2013	Carnegie Mellon University Herbert A. Simon Award for Teaching Excellence in Computer Science
2012	Carnegie Mellon University Allen Newell Award for Research Excellence.
2011	Sloan Foundation Fellowship
2010–2012	Winner, JouleSort competition for Energy-Efficient Sorting.
2009	Best paper award, Symposium on Operating Systems Principles (SOSP)

2006–2008	Selected to serve on the DARPA Computer Science Study Panel
2006	NSF CAREER Award (Faculty Early Career Development)
2005	MIT EECS George M. Sprowls Award for outstanding Ph.D. thesis
2002–2004	Microsoft Research Graduate Fellowship
2001	Best Student Paper, 8th IEEE Workshop on Hot Topics in Operating Systems
2001	MIT Joseph Levin award for best MasterWorks oral presentation
1999	MIT Vinton Hayes Fellowship (graduate)
1998	University of Utah Graduating Student Leadership Award
1993	Member, Phi Kappa Phi and Golden Key academic honor societies
1993–1997	University of Utah Honors at Entrance Scholarship
1993	National Merit Scholar

Service and Other Activities

2016	Program Committee, ASPLOS.
2010, 2014	Program Committee, OSDI 2010
2014	Program Committee, Symposium on Experimental Algorithms
2010–2013	Member, DARPA Information Sciences and Technology (ISAT) advisory group.
2006–2014	NSF panelist/proposal reviewer, multiple panels.
2012, 2006	Program Committee, Network Systems Design and Implementation (NSDI) 2006.
2012, 2009	Program committee, SOSP 2013 (“Heavy”)
2011	Program co-chair, NSDI.
2010	Participant, NSF Future Directions for Computer Systems Research workshop.
2010	Organizing committee, 2nd workshop on Architectural Concerns in Large Datacenters, in conjunction with ISCA 2010.
2010	Program Committee, 1st International Conference on Energy-Efficient Computing and Networking
2008, 2009	Program Committee, SIGCOMM 2009 (“Heavy”)
2008	Program co-chair, Workshop on Hot Topics in Networking (HotNets)
2007, 2011	Program Committee, Workshop on Hot Topics in Networking (HotNets)
2006, 2010–11	Consulting: Intel Research/Labs, Pittsburgh.
2003–2006	Consulting: Banner & Witcoff, attorneys at law.
2006	Program Committee, Internet Measurement Conference.
2006	Program co-chair, WORLDS 2006.
2006	Program Committee, 2nd Workshop on Hot Topics in Systems Dependability (HotDep).
2005	Editor (one of twelve), “Report of the NSF Workshop on Research Challenges in Distributed Computer Systems”
2005	Program Committee and Works-in-progress chair, USENIX 2005
2004–2005	Program Committee, Workshop on Real, Large, Distributed Systems (WORLDS).
	Reviewer for OSDI, SOSP, SIGCOMM, NSDI, CCR, HotOS, ToN, IEEE TDSC, Infocom, HotNets.
1999–2003	Secretary, board member, and rock climbing instructor for the MIT Outing Club.

1999–2000 Secretary, Utah Regional Exchange Point
Member, IEEE, ACM, USENIX.

Research - Energy Efficient Data-Intensive Computing

- 2012–present **Memory-Efficient, Architecture-Aware Data-Intensive Computing** CMU
- In my ongoing work, my students, collaborators, and I are designing techniques for managing and computing on high volumes of data that will work well with the computer architectures of the future: Higher core counts, the requirement for more parallelism, the need to exploit vectorization and SIMD, and the availability of increasingly ASIC-specialized computation for individual domains such as machine learning.
- 2007–2012 **FAWN: A Fast Array of Wimpy Nodes** CMU
- Through the FAWN project, I am exploring the design of highly energy efficient clusters for data-intensive computing. FAWN constructs clusters from large numbers of relatively “wimpy” embedded systems. It exploits fundamental efficiencies of using slower processors, and is designing algorithmic and systems techniques to mask the complexity of programming and managing systems that operate at increased scale with decreased per-node capability.
- 2007–2010 **Network Support for Data-Intensive Computing** CMU
- This project explores the networking challenges that arise in modern datacenters and devises solutions to these challenges. Examples include our work on overcoming datacenter “incast”, a problem in which the network can drastically slow down when tens or thousands of computers all communicate with one receiver.

Research - Network Architecture, Analysis, and Resilience

- 2010– **The eXtensible Internet Architecture (XIA)** CMU
- I am one of the principal investigators on a multi-year, multi-million dollar NSF-funded effort to develop the fundamental underpinnings of a new, more flexible, and more secure Internet Architecture.
- 2007–2010 **The Accountable Internet Protocol** CMU
- Together with my collaborators at MIT, Berkeley, and Georgia Tech, I am developing a novel framework for building a more secure Internet. The AIP project is based upon the notion of using self-certifying addresses instead of IP addresses (a self-certifying address is the hash of a public key). We have thus far shown that using this foundation can greatly simplify many aspects of providing network security, including reducing the potential for Denial-of-Service attacks and enabling simpler, self-configuring secure routing.
- Under AIP, we have also explored pragmatic alternatives to conventional routing security. Instead of using central authorities to cryptographically authenticate routing information, we explored the use of purely end-to-end authentication (which AIP facilitates) together with multi-path routing, showing that this approach is more robust to route hijacking than conventional approaches such as S-BGP, without requiring cryptographic authentication of routing announcements.
- 2006– **Perspectives** CMU
- Perspectives takes ideas from overlays and multi-path networking to and applies them to authenticating remote computers. The system has two goals: First, to materially improve network (particularly Web) security for ordinary users by enabling the easy and safe use of self-signed certificates. Second, to explore the utility of creating an “automatic” public key infrastructure based upon long-term observations. Perspectives is currently available as a Firefox browser plugin and as a patch to SSH. These plugins use a simple method to authenticate a self-signed certificate received from a server: They contact a set of “notary” servers scattered around the network. These servers inform the client what key they observe the server using, and for how long they have observed it. As a result, for an attacker to successfully deceive the client into accepting a false certificate, the attacker must have controlled all paths to the server for a long period of time. The Firefox plugin has been downloaded

by over 30,000 users, and the availability of the technique has presented a new answer to the debate about how browsers should handle self-signed certificates.

2005–

Data-Oriented Transfer

CMU

The Data-Oriented Transfer project is exploring a new architecture for applications that perform bulk data transfers. This architecture, called DOT (for *data-oriented transfer*), cleanly separates two functions that are co-mingled in today's applications. Using DOT, applications perform content *negotiation* to determine what content to send. They then pass that data object to the transfer service to perform the actual data transmission. This separation increases application flexibility, enables the rapid development of innovative transfer mechanisms, reduces developer effort, and allows increased efficiency through cross-application sharing of cached data.

In addition to the core architecture, the DOT project has developed a number of new transfer techniques. *SET*, or Similarity-Enhanced Transfer, is a peer-to-peer system that uses a scalable algorithmic approach to locate not only sources of the exact file a client wishes to download, but also similar copies, such as a truncated or slightly modified version. *Dsync* is a file synchronization tool that provides the benefits of two-node file synchronization tools such as rsync and the multiple-node efficiency of peer-to-peer transfers.

Finally, we have examined extensively the use of content addressability to increase the efficiency and robustness of networks, particularly wireless networks. *RTS-id* is a simple, fully backwards-compatible addition to 802.11 wireless that enables nodes to suppress transmissions of packets they have overheard (e.g., in a previous hop in a multi-hop mesh network). *Ditto* uses DOT's content-centric transfers to allow wireless mesh nodes to cache data that they overhear being transferred between other nodes. In scenarios where all clients eventually want the same data, such as disseminating popular software upgrades, Ditto can improve mesh network throughput by up to 10x.

2005–

The Datapository

CMU

The datapository is a shared network measurement storage and analysis infrastructure, designed to unite network data collection and analysis efforts at CMU and elsewhere. The datapository consists of a set of hardware resources (storage and computation), along with schema definitions, standard interfaces for analysis tools, and a set of tools for manipulating stored data (e.g., end-to-end probing data, routing information, topology snapshots). We are building the datapository in collaboration with researchers at MIT, Georgia Tech, and the University of Utah's network experimentation testbed, Emulab.

1999–2005

Resilient Overlay Networks and MONET

MIT

My dissertation research investigated host-based techniques that improve the end-to-end fault resilience of communication on the Internet. RON creates dynamic overlay networks between participating hosts or applications. The overlay networks use a combination of active probing and passive measurements to find more reliable and better performing routes by sending packets through the other participating nodes in the overlay. Results from this research showed that RON-like approaches can avoid up to half of the failures that interrupt communication and can significantly improve latency for poorly-performing paths. A set of Internet-based experiments in 2001 showed that RON can avoid up to half of the failures that interrupt communication, and can offer significant latency improvements for poorly-performing paths. MONET extends this by including multiple physical paths from sites and by moving from a host-centric view to a server-centric view (in which clients could be connected to one of several server replicas). MONET's combination of techniques can improve availability by an order of magnitude compared to current approaches such as BGP multi-homing.

1999–2000

Congestion Manager

MIT

The Congestion Manager provides a unified congestion controller for ensembles of TCP and UDP flows that eliminates adverse interactions and extends the benefits of congestion control to non-TCP applications. To help evaluate the CM, I co-implemented a congestion-controlled version of *vat*, an internet audio tool, which used the Congestion Manager to behave in a TCP-friendly manner with

low overhead. I helped design and implement the kernel to user API for the CM, and performed extensive performance measurements of the CM for both in-kernel and userspace applications.

1998–

Emulab + RON Testbed

University of Utah / MIT

Systems and networking researchers frequently use home-grown testbeds to evaluate prototypes and perform Internet measurements. To reduce the burden of creating these testbeds and to help provide a framework with better experimental repeatability, I played a part in the conception and design of a large-scale network testbed, Emulab, and a portion of its management databases, algorithms, and software. I deployed (and still manage) the 36-node RON Internet testbed, one of the first successful “overlay” network testbeds.

Research - Network Security

2003

Mayday: Distributed Filtering for Internet Services

MIT

Mayday presents an incrementally deployable Denial of Service *prevention* service that acts primarily as an overlay service, minimizing the network changes required for its deployment. Unlike tactics such as spoofing prevention, Mayday provides immediate protection to its deployers instead of requiring upgrades on the part of third parties. Mayday generalizes earlier work on Secure Overlay Services by separating overlay routing from filtering and by providing a larger set of choices for each, allowing the implementer to choose a high-performance deployment such as proximity routing, or a slower system that can withstand more capable attackers.

As part of the evaluation of Mayday and earlier work, I developed several practical attacks, two of them novel, that are effective against filtering-based systems like Mayday and SOS.

Summer 2001

Secure Network Attached Disks

Compaq SRC

Traditional disk architectures interpose a fileserver between clients and disks to provide access control. *Network Attached Disk* efforts aim to place the disks directly on the network, eliminating the bottleneck presented by the file server. The capability-based approach we examined permits the disks to export a familiar block-based interface; compared to earlier NAD efforts, this eliminates disk layout changes and simplifies the on-disk implementation. I created a filesystem simulator for our proposed architecture and created a benchmark suite from measurements of SRC's fileserver traffic to drive the simulator.

1997-1999

Flask: A secure microkernel

University of Utah

Users' requirements for operating systems vary considerably, from the MLS policies favored in military applications, to RBAC-like policies more common in large enterprises, to type enforcement policies favored for providing least privilege to local processes. The Flask security architecture provides fine-grained access rights and permits for their revocation to permit a single OS implementation to support a wide range of security policies. As an undergraduate, and continuing as research staff, I implemented and benchmarked parts of the Flask architecture, improved the reliability of the underlying Fluke microkernel, and implemented several of the example applications used in its evaluation. The technology developed for Flask later became an integral component of the SELinux secure operating system.