Jeffrey C. Mogul

Curriculum Vitae

Aug 2020

Contact 320 O'Connor Street, Menlo Park, CA 94025

jeffmogul@acm.org

Research Interests

Computer networking, especially Software-Defined Networks, Data-Center networks, and Cloud networks.

Operating systems and their interaction with computer architecture

Employment

5/2013–Present	Principal Software Engineer, Network Infrastructure Google, LLC.
5/2002–4/2013	HP Fellow *HP Labs, Networking and Communications Lab (NCL)*. Research on Software-Defined Networks, data-center networks, and cloud networks. Previous HP Labs groups: *Western Research Lab; Internet Systems and Storage Lab; Enterprise Systems and Software Lab; Exascale Computing Lab*. Research on computer systems design and implementation, including computer networking, operating systems, and computer architecture.
8/2000–5/2002	Compaq Staff Fellow Compaq Computer Corporation Western Research Laboratory.
2/2000–8/2000	Principal Member of Technical Staff Compaq Computer Corporation Western Research Laboratory.
6/1998–2/2000	Senior Member of Technical Staff Compaq Computer Corporation Western Research Laboratory.
9/1992–6/1998	Senior Member of Technical Staff Digital Equipment Corporation Western Research Laboratory.
7/1986–6/1998	Principal Software Engineer Digital Equipment Corporation Western Research Laboratory.
9/1981–2/1986	Research Assistant Computer Science Department, Stanford University.
9/1980–9/1981	Systems Programmer Computer Systems Laboratory, Stanford University.

Education

Ph.D., Computer Science, March 1986
Stanford University, Department of Computer Science
Thesis: *Representing Information About Files*
Adviser: Prof. Brian Reid
Committee: Prof. Forest Baskett and Prof. John L. Hennessy

M.S., Computer Science, October 1980
Stanford University, Department of Computer Science
Masters project: *A U-code to VAX code translator*
Adviser: Prof. John L. Hennessy

S.B., Computer Science, June 1979
Massachusetts Institute of Technology,
Department of Computer Science and Electrical Engineering
Thesis: *Attaching the VAX-11/780 to the 'Local Network'*
Adviser: Prof. Stephen A. Ward

Honors

Fellow, Association for Computing Machinery
Awarded in 2002

Recent Projects at HP Labs (selected)

Modular composition of SDN controllers: The long-term success of Software-Defined Networking (SDN) depends on rapid deployment of new network features using flexible controller platforms, often referred to as the "app-store model." This model requires a platform to support modular composition of applications, just as the OpenFlow protocol aims to support modular composition of controllers and switch data-planes.

Almost all prior work on SDN controllers has focussed on getting the control system to do what one wants it to do; very little work has addressed the problem of how to decide among competing choices in a resource-constrained network. I initiated and led the "Corybantic" project at HP Labs, which tries to dynamically maximize the net economic value of a network infrastructure. Modules reserve resources (such as link bandwidth) via proposals, and evaluate each others' proposals, under the control of a coordinator. Modules know nothing about each other except for the reservations, proposals, and evaluations communicated via the coordinator. We believe that this search-based approach is more likely to succeed than one based on numerical optimization.

Cloud-network bandwidth guarantees: Infrastructure-as-a-Service (IaaS) clouds typically provide some performance guarantees for CPU, memory, and storage, but not for networking. We expect many enterprise applications will require or benefit from bandwidth guarantees if they are to migrate to the cloud. Numerous researchers have proposed approaches for cloud-network bandwidth guarantees, but we still lack consensus on the best approach. I have led or participated in several projects aimed at clarifying these questions, including:

- Each research project seems to describe not only a unique approach to enforcing IaaS

network guarantees, but also a unique model for defining those guarantees. I wrote (with Lucian Popa) a survey article describing many such models, and suggesting some criteria for understanding which ones are most appropriate.

- One challenge is that cloud-tenant bandwidth requirements vary temporally (e.g., Wednesday noon vs. Saturday night) and spatially (some VM-to-VM paths are unused, others are heavily used.) "Cicada" shows that it is possible to predict, via machine learning, pairwise bandwidth requirements in the presence of both temporal and spatial variation. (With Prof. Hari Balakrishnan and Katrina LaCurts, of MIT)
- In a team lead by Lucian Popa, we have developed a hypervisor-based mechanism that supports both "hose-model" bandwidth guarantees (i.e., for the sum of a VM's pairwise connections) and work-conserving allocation of spare bandwidth.

SDN switch design: Based on work initiated by colleagues at HP Labs, who ported OpenFlow to HP's switch hardware, I initiated the DevoFlow design, which extends OpenFlow so as to greatly reduce the control traffic required to support complex, rapidly-responding traffic engineering and similar SDN applications. DevoFlow introduces "rule-cloning" and "triggers," relatively simple changes to the basic OpenFlow protocol

DevoFlow creates a dilemma: existing high-performance OpenFlow switches process counter updates on-ASIC, but this makes it difficult to introduce or modify the set of triggers. Also, OpenFlow generally creates scaling problems for on-ASIC counters. In response, I developed the "Software-Defined Counters" design, which exploits high-speed DRAM and relatively cheap embedded CPU cycles to move counter maintenance into software, allowing better scaling and much more flexibility.

Scalable data-center networks: Large data centers require large, high-performance, low-cost networks. Working on projects led primarily by Jayaram Mudigonda and Praveen Yalagandula, we developed several novel techniques:

- **SPAIN:** Most high-performance data-center networks require multipath topologies to achieve high bandwidth and resiliency with low-cost switches. SPAIN achieves multipath scaling using unmodified commodity switches, with end-host support, through the use of advanced heuristic algorithms to map the set of possible paths onto a small set of trees. These trees can then be instantiated as VLANs, supporting multiple edge-disjoint paths between pairs of end hosts.
- **NetLord:** Large networks, especially IaaS networks with multiple VMs per server, can expose hardware switches to large numbers of layer-2 (MAC) addresses. Cheap switches have small address-learning tables, which leads to pathological behavior with large sets of active addresses. NetLord uses encapsulation to hide these addresses from switches. NetLord slightly pre-dates the VXLAN protocol, and exploits existing MAC and IP headers, rather than requiring a new protocol standard.
- **Flying Cable Monster:** The use of multipath topologies composed of relatively inexpensive, and hence small, switches creates a huge variety of options: topology family (e.g., FatTree or HyperX), switch radix, number of downlinks per switch. Network designers need to pare down from billions of potential designs to a few good ones. We developed

heuristic algorithms to choose topologies that optimize bandwidth/$, including both capital and installation costs for cables. Our algorithms also try to optimize the physical placement of switches.

- **Scalable Multicast:** We expect some enterprise applications will require multicast support before they can migrate to an IaaS cloud. Existing clouds do not support multicast, probably because inexpensive switches have multicast tables that are too small for IaaS scales. We developed an algorithm that improves on the prior state of the art ("Dr. Multicast," from Cornell/IBM) to support significantly more multicast groups in a given infrastructure.

Most of these projects involved proposed or successful technology transfers to HP's businesses.

Co-design of memory hardware and operating system support: For many years, computer systems designers mostly took memory (DRAM) for granted, especially as it continually improved due to Moore's Law. DRAM now faces serious constraints for both power and density. While novel "storage-class memory" may some day resolve the problems with DRAM, we investigated several shorter-term approaches:

- **FLAM:** Others had proposed designing main memory as a hybrid of DRAM and flash; DRAM would hold rapidly-changing data, while flash could hold long-lived, read-mostly data. Similar hybrids could be constructed using phase-change memory or memristors, instead of flash. I initiated the FLAM project, which explored the possibility that the operating system could optimize the use of such hybrids, through its knowledge of various meta-data about individual pages. We provided some initial validation through simulation.
- **Energy-efficient DRAMs that exploit operating system knowledge:** Traditional DRAMs use large "row buffers" to capture the locality in memory references. This approach has worked well for single-core CPUs, but multi-core processing decreases short-term memory locality. The memory controller thus often fetches a large row buffer but uses only a small part, before discarding the rest; this wastes significant DRAM energy. Prior work by Naveen Muralimanohar had demonstrated that fetching smaller units could save energy, but it was not clear how to choose a fetch size that would be optimal across many applications. In our "Dynamic Multiple-Subarray Access" (DMSA) design, the operating system uses a modified TLB and memory controller (MC), and a simple feedback loop, to dynamically choose the right fetch size for each page. Initial simulations suggested some value for this approach.

  Similarly, memory designers have a choice between error-correcting codes (ECC) and error-detecting codes, such as parity. We proposed a simple change to the TLB and MC that would allow the operating system to choose between error correction and error detection on a per-page basis. For pages where the OS could reconstruct the data, error detection is sufficient and could save significant DRAM energy.

Older Projects (selected)

Asymmetric Single-ISA Multicore processors: The ratio between core size (and power) and core performance is not constant; small cores tend to have better performance per watt. Other

researchers had previously shown that many applications can be migrated between small cores and large cores to gain better throughput/watt, if all cores support the same Instruction-Set Architecture (ISA).

I initiated a project based on the observation that operating system code (more generally, "operating-system-like code") typically does not benefit as much from the higher performance of complex cores, nor does it often use features such as floating point hardware. We developed an "Asymmetric Single-ISA" (ASISA) design, in which the kernel rapidly switched between cores on entry to and exit from certain long-running system calls. Through simulations, we showed that ASISA could save energy. We also radically improved the speed of core-to-core migration in an experimental Linux kernel.

Performance debugging for distributed applications: Large distributed systems are often hard to debug, and it can be even harder to localize performance bugs in such systems. I led a team that developed a series of experimental systems to test several approaches:

- **Project5:** In this work, we assumed that the system components were black boxes – we could not look at or modify their source code – and we developed several different algorithms to analyze message traffic between components. Our analyses identified the causal structure of communication within a system, and attempted to localize the high-delay nodes on slow paths.
- **Wide-Area Project5:** This work extended Project5 to wide-area networks, in which network delays are significant.
- **Pip:** With Pip, we abandoned the black-box model, and instead assumed that programmers could annotate their code with a small number of measurement points, and that they could write "expectations" (in a simple language) about how they believed their system would perform. Pip would then monitor system performance and reveal which expectations, if any, were being violated. Pip included a simple GUI and a tool to reverse-engineer expectations from a properly-running system.

Patrick Reynolds and Janet Wiener contributed significantly to this work, which became Patrick's dissertation.

HTTP protocol design and optimization: I contributed significantly to the design of HTTP/1.1, including the experimental validation of the "persistent connections" model, of pipelined requests, and of certain kinds of prefetching. I also worked on various aspects of caching in HTTP, including experimental studies and protocol design. Venkat Padmanabhan contributed significantly to some of this work.

Internet performance: I conducted or led several research projects aimed at improving Internet implementations, including studies of TCP "ACK compression" and reference locality, and of approaches to reduce or eliminate "Receive Livelock" (with KK Ramakrishnan).

Together with Guarav Banga, I studied the bottlenecks in Internet server performance. We developed several kernel changes to radically improve scaling of Internet servers, and we developed the concept of "Resource Containers," to allow a multi-user server to properly account for CPU resources.

UNIX-based Firewall: In response to the first widespread Internet worm, I developed the first easily-configurable UNIX-based firewall. The "screend" software was widely distributed as open source, and was the basis for an early firewall product from Digital Equipment Corp.

Spritely NFS: The NFS protocol does not directly provide cache consistency, which means that clients must either give up the performance benefits of aggressive caching, or tolerate unpredictable behavior when files are write-shared. In Spritely NFS, we grafted a consistency protocol from the Sprite operating system onto NFS, and showed that it could improve performance while maintaining consistency. We also developed crash-recovery mechanisms for this design.

Packet Filter: Mike Accetta and Rick Rashid developed the first version of a UNIX-based packet filter, which allowed user-level implementations of network protocol stacks, and also made it possible to develop packet-sniffing tools such as tcpdump. I generalized their design and made it more efficient, and introduced it into the BSD software distribution. It was ultimately supplanted by the improved "bpf" packet filter design.

Publications

Note: many of the URLs below are probably out of date.

Top-cited Papers

These papers have all been cited at least 300 times (per Google Scholar, 4 March 2013, in declining order of citation count): (These papers are also listed below.)

Journal Papers

Deborah Estrin, Jeffrey C. Mogul, and Gene Tsudik. Visa Protocols for Controlling Inter-Organization Datagram Flow. *IEEE Journal on Selected Areas in Communications*, 7(4):486–498, May 1989.

Jeffrey C. Mogul. Network Locality at the Scale of Processes. *ACM Transactions on Computer Systems*, 10(2):81–109, May 1992.

Jeffrey C. Mogul. Recovery in Spritely NFS. *Computing Systems*, 7(2):201–262, Spring 1994.

Venkata N. Padmanabhan and Jeffrey C. Mogul. Improving HTTP Latency. *Computer Networks and ISDN Systems*, 28:25–35, 1995. Slightly revised version of paper in Proc. 2nd International WWW Conf. '94: Mosaic and the Web.

Jeffrey C. Mogul and K. K. Ramakrishnan. Eliminating Receive Livelock in an Interrupt-driven Kernel. *ACM Trans. on Computer Systems*, 15(3):217–252, August 1997.

Jeffrey C. Mogul. Server-Directed Transcoding. *Computer Communications*, 24(2):155–162, February 2001.

Bjorn Knutsson, Honghui Lu, Jeffrey Mogul, and Bryan Hopkins. Architecture and Performance of Server-Directed Transcoding. *ACM Transactions on Internet Technology*, 3(4):392–424, November 2003.

Jeffrey C. Mogul. Clarifying the Fundamentals of HTTP. *Software–Practice and Experience*, 34(2):103–134, February 2004.

Richard Strong, Jayaram Mudigonda, Jeffrey C. Mogul, Nathan Binkert, and Dean Tullsen. Fast Switching of Threads Between Cores. *Operating Systems Review special issue on The Interaction Among the OS, the Compiler, and Multicore Processors*, 43(2):35–45, April 2009.

Books

Jeffrey C. Mogul. IP Network Performance. In Dan Lynch and Marshall Rose, editors, *Internet System Handbook*, chapter 15, pages 575–675. Addison-Wesley, Reading MA, 1992. Editors: Dan Lynch and Marshall Rose; ISBN 0201567415.

Refereed Conference Papers

J. Mogul. Representing information about files. In *Proc. ICDCS*, pages 432–439. IEEE, May 1984.

Christopher A. Kent and Jeffrey C. Mogul. Fragmentation Considered Harmful. In *Proc. SIGCOMM '87 Workshop*, pages 390–401, Stowe, VT, August 1987.

Jeffrey C. Mogul, Richard F. Rashid, and Michael J. Accetta. The Packet Filter: An Efficient Mechanism for User-Level Network Code. In *Proc. SOSP*, pages 39–51, Austin, Texas, November 1987.

David R. Boggs, Jeffrey C. Mogul, and Christopher A. Kent. Measured Capacity of an Ethernet: Myths and Reality. In *Proc. SIGCOMM '88 Symposium on Communications Architectures and Protocols*, pages 222–234, Stanford, CA, August 1988. ACM SIGCOMM.

Jeffrey C. Mogul. Simple and Flexible Datagram Access Controls for Unix-based Gateways. In *Proc. USENIX Summer*, pages 203–221, Baltimore, MD, June 1989.

V. Srinivasan and Jeffrey C. Mogul. Spritely NFS: Experiments with Cache-Consistency Protocols. In *Proc. SOSP*, pages 45–57, Litchfield Park, AZ, December 1989.

Jeffrey C. Mogul. Efficient Use Of Workstations for Passive Monitoring of Local Area Networks. In *Proc. SIGCOMM*, pages 253–263, Philadelphia, PA, September 1990. ACM SIGCOMM.

Jeffrey C. Mogul and Anita Borg. The Effect of Context Switches on Cache Performance. In *Proc. ASPLOS*, pages 75–84, Santa Clara, CA, April 1991.

Jeffrey C. Mogul. Network Locality at the Scale of Processes. In *Proc. SIGCOMM*, pages 273–284, Zurich, September 1991.

Jeffrey C. Mogul. Observing TCP Dynamics in Real Networks. In *Proc. SIGCOMM*, pages 305–317, Baltimore, MD, August 1992.

Jeffrey C. Mogul. A Better Update Policy. In *Proc. USENIX Summer*, pages 99–111, Boston, MA, June 1994. **Best Paper Award**.

Venkata N. Padmanabhan and Jeffrey C. Mogul. Improving HTTP Latency. In *Proc. 2nd International WWW Conf. '94: Mosaic and the Web*, pages 995–1005, Chicago, IL, October 1994. URL http://www.ncsa.uiuc.edu/SDG/IT94/Proceedings/DDay/mogul/HTTPLatency.html.

Jeffrey C. Mogul, Joel F. Bartlett, Robert N. Mayo, and Amitabh Srivastava. Performance Implications of Multiple Pointer Sizes. In *Proc. USENIX*, pages 187–200, New Orleans, LA, January 1995. **Best Paper Award**.

Jeffrey C. Mogul. The Case for Persistent-Connection HTTP. In *Proc. SIGCOMM*, pages 299–313, Cambridge, MA, August 1995.

Jeffrey C. Mogul and K. K. Ramakrishnan. Eliminating Receive Livelock in an Interrupt-driven Kernel. In *Proc. USENIX 1996 Technical Conference*, pages 99–111, San Diego, CA, January 1996. URL http://www.usenix.org/publications/library/proceedings/sd96/mogul.html.

Jeffrey C. Mogul, Fred Douglis, Anja Feldmann, and Balachander Krishnamurthy. Potential benefits of delta encoding and data compression for HTTP. In *Proc. SIGCOMM '97 Conference*, pages 181–194, Cannes, France, September 1997. ACM SIGCOMM.

Fred Douglis, Anja Feldmann, Balachander Krishnamurthy, and Jeffrey Mogul. Rate of Change and Other Metrics: a Live Study of the World Wide Web. In *Proc. Symposium*

on Internet Technologies and Systems, pages 147–158, Monterey, CA, December 1997. USENIX.

Thomas M. Kroeger, Darrell D. E. Long, and Jeffrey C. Mogul. Exploring the Bounds of Web Latency Reduction from Caching and Prefetching. In *Proc. Symposium on Internet Technologies and Systems*, pages 13–22, Monterey, CA, December 1997. USENIX. URL http://www.cse.ucsc.edu/

Gaurav Banga and Jeffrey C. Mogul. Scalable kernel performance for Internet servers under realistic loads. In *Proc. 1998 USENIX Annual Technical Conf.*, pages 1–12, New Orleans, LA, June 1998. USENIX. **Best Student Paper Award/Best Paper Award**.

Gaurav Banga, Peter Druschel, and Jeffrey C. Mogul. Resource Containers: A New Facility for Resource Management in Server Systems. In *Proc. 3rd. Symp. on Operating Systems Design and Implementation*, pages 45–58, New Orleans, LA, February 1999. **Best Student Paper Award**.

Balachander Krishnamurthy, Jeffrey C. Mogul, and David M. Kristol. Key Differences between HTTP/1.0 and HTTP/1.1. In *Proc. WWW8*, Toronto, May 1999. URL http://www.research.att.com/library/trs/TRs/98/98.39/98.39.1.body.ps.

Gaurav Banga, Jeffrey C. Mogul, and Peter Druschel. A scalable and explicit event delivery mechanism for UNIX. In *Proc. 1999 USENIX Technical Conference*, pages 253–266, Monterey, CA, June 1999. **Outstanding Paper Award**.

Terence Kelly and Jeffrey Mogul. Aliasing on the World Wide Web: Prevalence and Performance Implications. In *Proc. 11th Intl. World Wide Web Conf.*, pages 281–292, Honolulu, HI, May 2002. URL http://www2002.org/CDROM/refereed/525.pdf.

Jeffrey C. Mogul. Clarifying the Fundamentals of HTTP. In *Proc. 11th Intl. World Wide Web Conf.*, pages 25–36, Honolulu, HI, May 2002. URL http://www2002.org/CDROM/refereed/444.pdf.

Bjorn Knutsson, Honghui Lu, and Jeffrey Mogul. Architecture and Pragmatics of Server-Directed Transcoding. In *Proc. 7th International Workshop on Web Content Caching and Distribution WCW-7*, pages 229–241, Boulder, CO, August 2002. URL http://2002.iwcw.org/papers/18500136.pdf.

Marcos K. Aguilera, Jeffrey C. Mogul, Janet L. Wiener, Patrick Reynolds, and Athicha Muthitacharoen. Performance Debugging for Distributed Systems of Black Boxes. In *Proc. 19th ACM Symposium on Operating Systems Principles SOSP-19*, pages 74–89, Bolton Landing, NY, October 2003. URL http://www.cs.rochester.edu/sosp2003/papers/p146-mogul.pdf.

Jeffrey C. Mogul, Yee Man Chan, and Terence Kelly. Design, Implementation, and Evaluation of Duplicate Transfer Detection in HTTP. In *Proc. First Symposium on Networked Systems Design and Implementation NSDI '04*, pages 43–46, San Francisco, CA, March 2004. URL http://www.usenix.org/events/nsdi04/tech/mogul.html.

Martin Arlitt, Balachander Krishnamurthy, , and Jeffrey C. Mogul. Predicting short-transfer latency from TCP arcana: A trace-based validation. In *Proc. Internet Measurement*

Conference, pages 213–226, Berkeley, CA, Oct 2005. URL http://www.imconf.net/imc-2005/papers/imc05efiles/arlitt/arlitt.pdf.

Patrick Reynolds, Charles Killian, Janet L. Wiener, Jeffrey C. Mogul, Mehul A. Shah, and Amin Vahdat. Pip: Detecting the unexpected in distributed systems. In *Proc. NSDI*, San Jose, CA, May 2006. URL `http://www.usenix.org/events/nsdi06/tech/full_papers/reynolds/reynolds.pdf`.

Patrick Reynolds, Janet L. Wiener, Jeffrey C. Mogul, Marcos K. Aguilera, Charles Killian, and Amin Vahdat. WAP5: Black-box Performance Debugging for Wide-Area Systems. In *Proc. WWW*, Edinburgh, May 2006. URL http://www2006.org/programme/files/pdf/2033.pdf.

Jeffrey C. Mogul. Emergent (mis)behavior vs. complex software systems. In *Proc. EuroSys 2006*, pages 293–304, Leuven, Belgium, April 2006. URL http://www.cs.kuleuven.ac.be/conference/EuroSys2006/papers/p293-mogul.pdf.

Jayaram Mudigonda, Praveen Yalagandula, Mohammad Al-Fares, and Jeffrey C. Mogul. SPAIN: COTS Data-Center Ethernet for Multipathing over Arbitrary Topologies. In *Proc. NSDI*, San Jose, CA, April 2010. URL `http://www.usenix.org/events/nsdi10/tech/full_papers/mudigonda.pdf`.

Erik Rubow, Rick McGeer, Jeffrey C. Mogul, and Amin Vahdat. Chimpp: A Click-based Programming and Simulation Environment for Reconfigurable Networking Hardware. In *Proc. ANCS*, La Jolla, CA, October 2010.

Jayaram Mudigonda, Praveen Yalagandula, Jeffrey C. Mogul, Bryan Stiekes, and Yanick Pouffary. NetLord: A Scalable Multi-Tenant Network Architecture for Virtualized Datacenters. In *Proc. SIGCOMM*, Aug. 2011.

Andrew R. Curtis, Jeffrey C. Mogul, Jean Tourrilhes, Praveen Yalagandula, Puneet Sharma, and Sujata Banerjee. DevoFlow: Scaling Flow Management for High-Performance Networks. In *Proc. SIGCOMM*, pages 254–265, Aug. 2011.

Jayaram Mudigonda, Praveen Yalagandula, and Jeffrey C. Mogul. Taming the Flying Cable Monster: A Topology Design and Optimization Framework for Data-center Networks. In *Proc. USENIX ATC*, Portland, OR, June 2011.

Lucian Popa, Praveen Yalagandula, Sujata Banerjee, Jeffrey C. Mogul, Yoshio Turner, and Jose Renato Santos. ElasticSwitch: practical work-conserving bandwidth guarantees for cloud computing. In *Proc. SIGCOMM*, pages 351–362, Hong Kong, 2013.

Seyed Kaveh Fayazbakhsh, Luis Chiang, Vyas Sekar, Minlan Yu, and Jeffrey C. Mogul. Enforcing Network-Wide Policies in the Presence of Dynamic Middlebox Actions using FlowTags. In *Proc. NSDI*, pages 543–546, Seattle, WA, 2014.

Alvin AuYoung, Yadi Ma, Sujata Banerjee, Jeongkeun Lee, Puneet Sharma, Yoshio Turner, Chen Liang, and Jeffrey C. Mogul. Democratic Resolution of Resource Conflicts Between SDN Control Programs. In *Proc.CoNEXT 2014*, pages 391–402, Sydney, Australia, 2014.

Refereed Workshop Papers

Anita Borg and Jeffrey C. Mogul. Position statement. In *Proc. Workshop on Workstation Operating Systems*, Cambridge, MA, November 1987. IEEE Technical Committee on Operating Systems.

Jeffrey C. Mogul. What is the right amount of statelessness in a file server? In *Proc. 2nd Workshop on Workstation Operating Systems*, pages 82–85, Pacific Grove, CA, September 1989. IEEE Technical Committee on Operating Systems.

Jeffrey C. Mogul. A Recovery Protocol for Spritely NFS. In *Proc. Workshop on File Systems*, pages 93–109, Ann Arbor, MI, May 1992. USENIX.

Jeffrey C. Mogul. SPECMarks are Leading Us Astray. In *Proc. 3rd Workshop on Workstation Operating Systems*, pages 160–161, Key Biscayne, FL, April 1992. IEEE Technical Committee on Operating Systems.

Jeffrey C. Mogul. Big Memories on the Desktop. In *Proc. 4th Workshop on Workstation Operating Systems*, pages 110–115, Napa, CA, October 1993. IEEE Technical Committee on Operating Systems.

Jeffrey C. Mogul. Operating Systems Support for Busy Internet Servers. In *Proc. Fifth Workshop on Hot Topics in Operating Systems HotOS-V*, page addendum, Orcas Island, Washington, May 1995. IEEE TCOS. Distributed but not bound with proceedings.

Jeffrey C. Mogul. Hinted caching in the Web. In *Proc. Seventh ACM SIGOPS European Workshop*, pages 103–108, Connemara, Ireland, September 1996. URL http://www-sor.inria.fr/sigops96/papers/mogul.ps.

Gaurav Banga, Peter Druschel, and Jeffrey C. Mogul. Better operating system features for faster network servers. In *Proc. Workshop on Internet Server Performance*, pages 69–79, Madison, WI, June 1998. URL http://www.cs.rice.edu/ gaurav/papers/wisp98.ps.

Jeffrey C. Mogul. Brittle metrics in operating systems research. In *Proc. HotOS-VII*, pages 90–95, Rio Rico, AZ, March 1999. URL http://www.research.digital.com/wrl/people/mogul/mogulhotos7.ps.

Greg Minshall, Yasushi Saito, Jeffrey C. Mogul, and Ben Verghese. Application Performance Pitfalls and TCP's Nagle Algorithm. In *Proc. 2nd Workshop on Internet Server Performance*, Atlanta, GA, May 1999. URL http://www.cc.gatech.edu/fac/Ellen.Zegura/wisp99/papers/minshall.ps.

Jeffrey C. Mogul. Server-Directed Transcoding. In *Proc. 5th International Web Caching and Content Delivery Workshop*, Lisbon, Portugal, May 2000. URL http://www.terena.nl/conf/wcw/Proceedings/S1/S1-4.ps.

Jeffrey C. Mogul. Position Summary: Toward a rigorous data type model for HTTP. In *Proc. 8th IEEE Workshop on Hot Topics in Operating Systems*, pages ??–??, Schloss Elmau, Germany, May 2001. IEEE TCOS. URL http://research.compaq.com/wrl/people/mogul/hotos8.

Jeffrey C. Mogul. TCP Offload Is a Dumb Idea Whose Time Has Come. In *Proc. 9th Workshop on Hot Topics in Operating Systems HotOS IX*, Lihue, HI, May 2003. URL http://www.usenix.org/events/hotos03/tech/mogul.html.

Jeffrey Mogul, Lawrence Brakmo, David E. Lowell, Dinesh Subhraveti, and Justin Moore. Unveiling the Transport API. In *Proc. 2nd Workshop on Hot Topics in Networks*, Cambridge, MA, November 2003.

Mema Roussopoulos, Mary Baker, David S. H. Rosenthal, TJ Giuli, Petros Maniatis, and Jeff Mogul. 2 P2P or Not 2 P2P? In *Proc. 3rd International Workshop on Peer-to-Peer Systems IPTPS '04*, San Diego, CA, February 2004. URL http://iptps04.cs.ucsd.edu/papers/roussopoulos-to-or-not.pdf.

John Wilkes, Jeffrey Mogul, and Jaap Suermondt. Utilification. In *Proc. 11th SIGOPS European Workshop*, Leuven, Belgium, September 2004. URL http://www.hpl.hp.com/research/ssp/papers/Utilification-final.pdf.

Jeffrey C. Mogul. Operating Systems Should Support Business Change. In *Proc. 10th Workshop on Hot Topics in Operating Systems HotOS X*, pages 43–48, Santa Fe, NM, June 2005. URL `http://www.usenix.org/events/hotos05/final_papers/mogul.html`.

Jeffrey C. Mogul and Martin Arlitt. SC2D: An Alternative to Trace Anonymization. In *Proc. MineNet-06*, pages 323–328, Pisa, Italy, September 2006. URL http://www.sigcomm.org/sigcomm2006/papers/minenet-08.pdf.

Mehul A. Shah, Mary Baker, Jeffrey C. Mogul, and Ram Swaminathan. Auditing to Keep Online Storage Services Honest. In *Proc. HotOS-XI*, San Diego, CA, May 2007. URL `http://www.usenix.org/event/hotos07/tech/full_papers/shah/shah.pdf`.

Jeffrey C. Mogul, Praveen Yalagandula, Jean Tourrilhes, Rick McGeer, Sujata Banerjee, Tim Connors, and Puneet Sharma. API Design Challenges for Open Router Platforms on Proprietary Hardware. In *Proc. HotNets-VII*, Oct. 2008. URL http://conferences.sigcomm.org/hotnets/2008/papers/hotnets08finalmogul.pdf.

Jeffrey C. Mogul, Eduardo Argollo, Mehul Shah, and Paolo Faraboschi. Operating System Support for NVM+DRAM Hybrid Main Memory. In *Proc. HotOS-XII*, May 2009. URL `http://www.usenix.org/events/hotos09/tech/full_papers/mogul/mogul.pdf`.

Erik Rubow, Rick McGeer, Jeff Mogul, and Amin Vahdat. Chimpp: A Modular Programming and Simulation Environment for Reconfigurable Networking Hardware. In *Proc. Workshop on Infrastructures for Software/Hardware co-design*, May 2010.

Jeffrey C. Mogul, Jean Tourrilhes, Praveen Yalagandula, Puneet Sharma, Sujata Banerjee, and Andrew R. Curtis. DevoFlow: Cost-Effective Flow Management for High Performance Enterprise Networks. In *Proc. Ninth ACM Workshop on Hot Topics in Networks HotNets-IX*, October 2010.

Jeffrey C. Mogul, Andrew Baumann, Timothy Roscoe, and Livio Soares. Mind the Gap: Reconnecting Architecture and OS Research. In *Proc. 13th Workshop on Hot Topics in Operating Systems (HotOS-XIII)*, May 2011.

Jeffrey C. Mogul and Paul Congdon. Hey, You Darned Counters! Get Off my ASIC! In *Proc. SIGCOMM Workshop on Hot Topics in Software Defined Networking (HotSDN)*, Aug. 2012.

Seyed Kaveh Fayazbakhsh, Vyas Sekar, Minlan Yu, and Jeffrey C. Mogul. FlowTags: enforcing network-wide policies in the presence of dynamic middlebox actions. In *Proceedings of the Second ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking, HotSDN 2013*, pages 19–24, Hong Kong, 2013.

Jeffrey C. Mogul, Jayaram Mudigonda, Jose Renato Santos, and Yoshio Turner. The NIC Is the Hypervisor: Bare-Metal Guests in IaaS Clouds. In *Proc. HotOS*, Santa Ana Pueblo, New Mexico, 2013.

Jeffrey C. Mogul, Alvin AuYoung, Sujata Banerjee, Lucian Popa, Jeongkeun Lee, Jayaram Mudigonda, Puneet Sharma, and Yoshio Turner. Corybantic: Towards the Modular Composition of SDN Control Programs. In *Proc. HotNets*, College Park, MD, 2013.

Seyed Kaveh Fayazbakhsh, Luis Chiang, Vyas Sekar, Minlan Yu, and Jeffrey C. Mogul. Extending SDN to Handle Dynamic Middlebox Actions via FlowTags. In *Open Networking Summit 2014 - Research Track, ONS 2014, Santa Clara, CA, USA, March 2-4, 2014*, 2014.

Katrina LaCurts, Jeffrey C. Mogul, Hari Balakrishnan, and Yoshio Turner. Cicada: Introducing Predictive Guarantees for Cloud Networks. In *Proc. HotCloud*, Philadelphia, PA., 2014.

Technical Reports

Jeffrey C. Mogul. Representing Information About Files. Technical Report STAN-CS-86-1103, Stanford University Computer Science Department, March 1986. (PhD Dissertation).

Jeffrey C. Mogul. The Leaf File Access Protocol: History and Specification. Technical Report STAN-CS-86-1137, Stanford University Computer Science Department, December 1986.

Jeffrey C. Mogul, Richard F. Rashid, and Michael J. Accetta. The Packet Filter: An Efficient Mechanism for User-Level Network Code. Research Report 87/2, Digital Equipment Corporation Western Research Laboratory, November 1987. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-87-2.html.

Christopher A. Kent and Jeffrey C. Mogul. Fragmentation Considered Harmful. Research Report 87/3, Digital Equipment Corporation Western Research Laboratory, November 1987. Expanded version of paper presented at SIGCOMM '87. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-87-3.html.

Jeffrey C. Mogul. The Experimental Literature of The Internet: An Annotated Bibliography. Research Report 88/3, Digital Equipment Corporation Western Research Laboratory, August 1988. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-88-3.html.

David R. Boggs, Jeffrey C. Mogul, and Christopher A. Kent. Measured Capacity of an Ethernet: Myths and Reality. Research Report 88/4, Digital Equipment Corporation Western Research Laboratory, September 1988. Expanded version of paper presented at SIGCOMM '88. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-88-4.html.

Deborah Estrin, Jeffrey C. Mogul, Gene Tsudik, and Kamaljit Anand. Visa Protocols for Controlling Inter-Organization Datagram Flow: Extended Description. Research Report

88/5, Digital Equipment Corporation Western Research Laboratory, December 1988. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-88-5.html.

Jeffrey C. Mogul. Simple and Flexible Datagram Access Controls for Unix-based Gateways. Research Report 89/4, Digital Equipment Corporation Western Research Laboratory, March 1989. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-89-4.html.

V. Srinivasan and Jeffrey C. Mogul. Spritely NFS: Experiments with and Implementation of Cache-Consistency Protocols. Research Report 89/5, Digital Equipment Corporation Western Research Laboratory, March 1989. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-89-5.html.

Jeffrey C. Mogul. Efficient Use of Workstations for Passive Monitoring of Local Area Networks. Research Report 90/5, Digital Equipment Corporation Western Research Laboratory, July 1990. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-90-5.html.

Jeffrey C. Mogul. Network Locality at the Scale of Processes. Research Report 91/11, Digital Equipment Corporation Western Research Laboratory, November 1991. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-91-11.html.

Jeffrey C. Mogul. Observing TCP Dynamics in Real Networks. Research Report 92/2, Digital Equipment Corporation Western Research Laboratory, April 1992. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-92-2.html.

Jeffrey C. Mogul. Recovery in Spritely NFS. Research Report 93/2, Digital Equipment Corporation Western Research Laboratory, June 1993. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-93-2.html.

Jeffrey C. Mogul. A Better Update Policy. Research Report 94/4, Digital Equipment Corporation Western Research Laboratory, April 1994. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-94-4.html.

Jeffrey C. Mogul, Joel F. Bartlett, Robert N. Mayo, and Amitabh Srivastava. Performance Implications of Multiple Pointer Sizes. Research Report 94/7, Digital Equipment Corporation Western Research Laboratory, December 1994. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-94-7.html.

Jeffrey C. Mogul. The Case for Persistent-Connection HTTP. Research Report 95/4, Digital Equipment Corporation Western Research Laboratory, May 1995. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-95-4.html.

Jeffrey C. Mogul. Network Behavior of a Busy Web Server and its Clients. Research Report 95/5, Digital Equipment Corporation Western Research Laboratory, October 1995. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-95-5.html.

Jeffrey C. Mogul and K. K. Ramakrishnan. Eliminating Receive Livelock in an Interrupt-driven Kernel. Research Report 95/8, Digital Equipment Corporation Western Research Laboratory, December 1995. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-95-8.html.

Jeffrey C. Mogul, Fred Douglis, Anja Feldmann, and Balachander Krishnamurthy. Potential benefits of delta encoding and data compression for HTTP. Research Report 97/4, Digital Equipment Corporation Western Research Laboratory, July 1997. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-97-4.html.

Gaurav Banga and Jeffrey C. Mogul. Scalable kernel performance for Internet servers under realistic loads. Research Report 98/6, Digital Equipment Corporation Western Research Laboratory, October 1998. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-98-6.html.

Jeffrey C. Mogul. A trace-based analysis of duplicate suppression in HTTP. Research Report 99/2, Compaq Western Research Laboratory, November 1999. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-99-2.html.

Jeffrey C. Mogul. Errors in timestamp-based HTTP header values. Research Report 99/3, Compaq Western Research Laboratory, December 1999. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-99-3.html.

Bill Hamburgen, Jeff Mogul, Brian Reid, Alan Eustace, Richard Swan, Mary Jo Doherty, and Joel Bartlett. Characterization of Organic Illumination Systems. Technical Note TN-13, Digital Equipment Corporation Western Research Laboratory, April 1989. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-TN-13.html.

Jeffrey C. Mogul and Anita Borg. The Effect of Context Switches on Cache Performance. Technical Note TN-16, Digital Equipment Corporation Western Research Laboratory, December 1990. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-TN-16.html.

Jeffrey C. Mogul. A Recovery Protocol for Spritely NFS. Technical Note TN-27, Digital Equipment Corporation Western Research Laboratory, April 1992. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-TN-27.html.

Jeffrey C. Mogul. Operating Systems Support for Busy Internet Servers. Technical Note TN-49, Digital Equipment Corporation Western Research Laboratory, May 1995. URL http://www.hpl.hp.com/techreports/Compaq-DEC/WRL-TN-49.html.

Jeffrey C. Mogul. Using screend to Implement IP/TCP Security Policies. Network Note NN-16, Digital Equipment Corporation Network Systems Laboratory, July 1991. Reissued as NSL Technical Note TN-2. URL http://www.hpl.hp.com/techreports/Compaq-DEC/NSL-TN-2.html.

Fred Douglis, Anja Feldmann, Balachander Krishnamurthy, and Jeffrey Mogul. Rate of Change and other Metrics: a Live Study of the World Wide Web. Technical Report 97.24.2, AT&T Labs – Research, December 1997. URL http://www.research.att.com/library/trs/TRs/97/97.24/97.24.2.body.ps.

Jeffrey C. Mogul, Yee Man Chan, and Terence Kelly. Design, Implementation, and Evaluation of Duplicate Transfer Detection in HTTP. Technical Report HPL-2004-29, HP Labs, February 2004. URL http://www.hpl.hp.com/techreports/2004/HPL-2004-29.html.

Martin Arlitt, Balachander Krishnamurthy, and Jeffrey C. Mogul. Predicting short-transfer latency from TCP arcana: extended version. Tech. Report HPL-2005-137, HP Labs, Sep 2005. URL http://www.hpl.hp.com/techreports/2005/HPL-2005-137.html.

Jeffrey C. Mogul. Operating Systems Should Support Business Change. Tech. Report HPL-2005-206, HP Labs, Nov 2005. URL http://www.hpl.hp.com/techreports/2005/HPL-2005-206.html.

Jeffrey C. Mogul. Emergent (Mis)behavior vs. Complex Software Systems. Tech. Report HPL-2006-2, HP Labs, Jan 2006. URL http://www.hpl.hp.com/techreports/2006/HPL-2006-2.html.

Jeffrey C. Mogul, Jarayam Mudigonda, Nathan Binkert, Partha Ranganathan, and Vanish Talwar. Operating Systems and Asymmetric Single-ISA CMPs: The Potential for Saving Energy. Tech. Report HPL-2007-140, HP Labs, Aug 2007. URL http://www.hpl.hp.com/techreports/2007/HPL-2007-140.html.

Jeffrey C. Mogul, Praveen Yalagandula, Jean Tourrilhes, Rick McGeer, Sujata Banerjee, Tim Connors, and Puneet Sharma. Orphal: API Design Challenges for Open Router Platforms on Proprietary Hardware. Tech. Report HPL-2008-108, HP Labs, Sep 2008. URL http://www.hpl.hp.com/techreports/2008/HPL-2008-108.html.

Rick Strong, Jayaram Mudigonda, Jeffrey Mogul, and Nathan Binkert. Initial Experiments in Visualizing Fine-Grained Execution of Parallel Software Through Cycle-Level Simulation. Tech. Report HPL-2008-210, HP Labs, Dec 2008. URL http://www.hpl.hp.com/techreports/2008/HPL-2008-210.html.

Jeffrey Mogul. IPv6 Node Discovery: Challenges and Approaches. Tech. Report HPL-2009-257, HP Labs, Sep 2009. URL http://www.hpl.hp.com/techreports/2009/HPL-2009-257.html.

Jayaram Mudigonda, Praveen Yalagandula, Mohammad Al-Fares, and Jeffrey C. Mogul. SPAIN: Design and Algorithms for Constructing Large Data-Center Ethernets from Commodity Switches. Tech. Report HPL-2009-241, HP Labs, Oct 2009. URL http://www.hpl.hp.com/techreports/2009/HPL-2009-241.html.

Jeffrey C. Mogul, Eduardo Argollo, Mehul Shah, and Paolo Faraboschi. Operating System Support for NVM+DRAM Hybrid Main Memory. Tech. Report HPL-2009-256, HP Labs, Nov 2009. URL http://www.hpl.hp.com/techreports/2009/HPL-2009-256.html.

Erik Rubow, Rick McGeer, Jeffrey C. Mogul, and Amin Vahdat. Chimpp: A Click-based Programming and Simulation Environment for Reconfigurable Networking Hardware. Tech. Report HPL-2010-25, HP Labs, Mar 2010. URL http://www.hpl.hp.com/techreports/2010/HPL-2010-25.html.

Jayaram Mudigonda, Praveen Yalagandula, and Jeffrey C. Mogul. Taming the Flying Cable Monster: A Topology Design and Optimization Framework for Data-center Networks. Tech. Report HPL-2011-30, HP Labs, Sep 2011. URL http://www.hpl.hp.com/techreports/2011/HPL-2011-30.html.

Lucian Popa, Praveen Yalagandula, Sujata Banerjee, Jeff Mogul, Yoshio Turner, and Jose Renato Santos. FluidCloud: Lubricating Cloud Network Performance. Tech. Report HPL-

2012-169, HP Labs, ??? 2012. URL http://www.hpl.hp.com/techreports/2012/HPL-2012-169.html.

Alvin AuYoung, Sujata Banerjee, Jeongkeun Lee, Jeffrey C. Mogul, Lucian Mudigonda, Jayaram; Popa, Puneet Sharma, and Yoshio Turner. Corybantic: Towards the Modular Composition of SDN Control Programs. Tech. Report HPL-2013-62, HP Labs, Sept 2013. URL http://www.hpl.hp.com/techreports/2013/HPL-2013-62.html.

Rachit Agarwal, Jayaram Mudigonda, Praveen Yalagandula, and Jeffrey C. Mogul. An Algorithmic Approach to Datacenter Cabling. Tech. Report HPL-2015-8, HP Labs, Feb 2015. URL http://www.hpl.hp.com/techreports/2015/HPL-2015-8.html.

Unrefereed Technical Papers

Jeffrey C. Mogul. The Experimental Literature of the Internet: An Annotated Bibiliography. *Computer Communication Review*, 19(1):58–71, January 1989. Slightly revised version of DECWRL Research Report 88/3.

Venkata N. Padmanabhan and Jeffrey C. Mogul. Using Predictive Prefetching to Improve World Wide Web Latency. *Computer Communication Review*, 26(3):22–36, 1996. URL http://daedalus.cs.berkeley.edu/publications/ccr-july96.ps.gz.

Gaurav Banga, Peter Druschel, and Jeffrey C. Mogul. Better operating system features for faster network servers. *Performance Evaluation Review*, 26(3):23–30, December 1998. URL http://www.cs.rice.edu/ gaurav/papers/WISP-PER.ps.

Jeffrey C. Mogul and Greg Minshall. Rethinking the TCP Nagle Algorithm. *Computer Communication Review*, 31(6):6–20, January 2001. URL http://www.acm.org/sigcomm/ccr/archive/2001/jan01/ccr-200101-mogul.html.

Vimalkumar Jeyakumar, Abdul Kabbani, Jeffrey C. Mogul, and Amin Vahdat. Flexible Network Bandwidth and Latency Provisioning in the Datacenter. *CoRR*, abs/1405.0631, 2014.

Editorial Papers

Jeffrey C. Mogul. Policies for the SIGOPS Hall of Fame Award. *Operating Systems Review*, 42(3):132–135, April 2008. URL http:???

Jeffrey C. Mogul. WOWCS '08: Workshop on Organizing Workshops, Conferences, and Symposia for Computer Systems. *;login:*, 33(4), August 2008. URL http://www.usenix.org/publications/login/2008-08/openpdfs/WOWCS08reports.pdf.

Jeffrey C. Mogul and Tom Anderson. Open Issues in Organizing Computer Systems Conferences. *SIGCOMM Comput. Commun. Rev.*, 38(3):93–102, July 2008.

Jeffrey C. Mogul. WOWCS: The Workshop on Organizing Workshops, Conferences, and Symposia for Computer Systems. *Operating Systems Review*, 43(2):106–107, April 2009.

Nathan Farrington, Nikhil Handigol, Christoph Mayer, Kok-Kiong Yap, and Jeffrey C. Mogul. Report on WREN 2009 – Workshop: Research on Enterprise Networking. *SIGCOMM Comput. Commun. Rev.*, 40(1):44–49, January 2010.

John W. Byers, Jeffrey C. Mogul, Fadel Adib, Jay Aikat, Danai Chasaki, Ming-Hung Chen, Marshini Chetty, Romain Fontugne, Vijay Gabale, László Gyarmati, Katrina LaCurts, Qi Liao, Marc Mendonca, Trang Cao Minh, S.H. Shah Newaz, Pawan Prakash, Yan Shvartzshnaider, Praveen Yalagandula, and Chun-Yu Yang. Report on the SIGCOMM 2011 conference. *SIGCOMM Comput. Commun. Rev.*, 42(1):80–96, January 2012.

Jeffrey C. Mogul. Towards more constructive reviewing of CS papers. In *Dagstuhl Perspectives Workshop on Publication Culture in Computing Research*, November 2012.

Jeffrey C. Mogul. Towards More Constructive Reviewing of SIGCOMM Papers. *Computer Communication Review*, 43(3):90–94, 2013.

Magazine Articles

Jeffrey C. Mogul. Fragmentation: Pros and Cons. *Connexions - The Interoperability Report*, 2(4):11–15, April 1988.

Jeffrey C. Mogul. Booting Diskless Hosts: The BOOTP Protocol. *Connexions - The Interoperability Report*, 2(10):14–18, October 1988.

Jeffrey C. Mogul. Subnetting: A Brief Guide. *Connexions - The Interoperability Report*, 3(1):2–9, January 1989.

Jeffrey C. Mogul. Path MTU Discovery. *Connexions - The Interoperability Report*, 4(12):12–15, December 1990.

Jeffrey C. Mogul. Applications and Techniques for LAN Monitoring. *Connexions - The Interoperability Report*, 5(10):36–45, October 1991.

Jeffrey C. Mogul. Speedier Squid: A Case Study of an Internet Server Performance Problem. *;login:*, 24(1):50–58, February 1999. URL http://www.usenix.org/publications/login/1999-2/squid.html.

Jeffrey C. Mogul. Squeezing More Bits Out of HTTP Caches. *IEEE Network*, 14(3):6–14, May/June 2000.

Submitted Papers

Jeffrey C. Mogul, Jayaram Mudigonda, Jose Renato Santos, and Yoshio Turner. The NIC is the Hypervisor: Bare-Metal Guests in IaaS Clouds. In *HotOS (in submission)*, 2013.

Lucian Popa, Praveen Yalagandula, Sujata Banerjee, Jeffrey C. Mogul, Yoshio Turner, and Jose Renato Santos. FluidCloud: Practical Work Conserving Bandwidth Guarantees for Clouds. In *SIGCOMM (anonymous submission)*, 2013.

Katrina LaCurts, Jeffrey C. Mogul, and Hari Balakrishnan. Cicada: Predictive Guarantees for Cloud Network Bandwidth. In *SIGCOMM (anonymous submission)*, 2013.

Standards

Ross Finlayson, Timothy Mann, Jeffrey Mogul, and Marvin Theimer. A Reverse Address Resolution Protocol. RFC 903, Network Information Center, June 1984. URL ftp://ds.internic.net/rfc/rfc903.txt.

Jeffrey Mogul. Internet Subnets. RFC 917, Network Information Center, October 1984. URL ftp://ds.internic.net/rfc/rfc917.txt.

Jeffrey Mogul. Broadcasting Internet Packets. RFC 919, NIC, October 1984. URL ftp://ds.internic.net/rfc/rfc919.txt.

Jeffrey Mogul. Broadcasting Internet Packets in the Presence of Subnets. RFC 922, Network Information Center, October 1984. URL ftp://ds.internic.net/rfc/rfc922.txt.

Jeffrey Mogul and Jon Postel. Internet Standard Subnetting Procedure. RFC 950, Network Information Center, August 1985. URL ftp://ds.internic.net/rfc/rfc950.txt.

Jeffrey C. Mogul, Christopher A. Kent, Craig Partridge, and Keith McCloghrie. IP MTU Discovery Options. RFC 1063, Network Information Center, July 1988. URL ftp://ds.internic.net/rfc/rfc1063.txt.

Jeffrey C. Mogul and Stephen Deering. Path MTU Discovery. RFC 1191, Network Information Center, November 1990. URL ftp://ds.internic.net/rfc/rfc1191.txt.

Jack McCann, Stephen E. Deering, and Jeffrey Mogul. Path MTU Discovery for IP version 6. RFC 1981, Network Information Center, August 1996. URL ftp://ds.internic.net/rfc/rfc1981.txt.

Roy T. Fielding, Jim Gettys, Jeffrey C. Mogul, Henrik Frystyk Nielsen, and Tim Berners-Lee. Hypertext Transfer Protocol – HTTP/1.1. RFC 2068, HTTP Working Group, January 1997.

Jeffrey C. Mogul, Roy T. Fielding, Jim Gettys, and Henrik Frystyk Nielsen. Use and Interpretation of HTTP Version Numbers. RFC 2145, HTTP Working Group, May 1997.

Jeffrey C. Mogul and Paul Leach. Simple Hit-Metering and Usage-Limiting for HTTP. RFC 2227, HTTP Working Group, October 1997.

Roy T. Fielding, Jim Gettys, Jeffrey C. Mogul, Henrik Frystyk Nielsen, Larry Masinter, Paul Leach, and Tim Berners-Lee. Hypertext Transfer Protocol – HTTP/1.1. RFC 2616, HTTP Working Group, June 1999.

Jeffrey C. Mogul, David L. Mills, Jan Brittenson, Jonathan Stone, and Ulrich Windl. Pulse-Per-Second API for UNIX-like Operating Systems, Version 1.0. RFC 2783, Network Working Group, March 2000.

Jeffrey C. Mogul and Arthur Van Hoff. Instance Digests in HTTP. RFC 3230, IETF, January 2002.

Jeffrey C. Mogul, Balachander Krishnamurthy, Fred Douglis, Anja Feldmann, Yaron Goland, Arthur van Hoff, and Daniel Hellerstein. Delta encoding in HTTP. RFC 3229, IETF, January 2002.

David G. Korn, Joshua P. MacDonald, Jeffrey C. Mogul, and Kiem-Phong Vo. The VCDIFF Generic Differencing and Compression Data Format. RFC 3284, IETF, June 2002.

Graham Klyne, Mark Nottingham, and Jeffrey C. Mogul. Registration Procedures for Message Header Fields. RFC 3864, IETF, September 2004.

Allyn Romanow, Jeffrey C. Mogul, Tom Talpey, and Stephen Bailey. Remote Direct Memory Access (RDMA) over IP Problem Statement. RFC 4297, IETF, December 2005.

Mark Nottingham and Jeffrey C. Mogul. HTTP Header Field Registrations. RFC 4229, IETF, December 2005.

Software

Jeffrey C. Mogul. XREF: Cross-Referencer for BASIC-PLUS Programs. DECUS RSTS-11-78, 1975. URL `http://pdp-11.trailing-edge.com/rsts11/rsts-11-078/index.html`.

Dave Curry and Jeff Mogul. NFSwatch. URL http://nfswatch.sourceforge.net/.

Bob Braden, Annette DeSchon, Jeff Mogul, and Amos Shapira. Nnstat 3.3beta, 1993. URL `ftp://ftp3.hp.com/ftp1/pub/alphaserver/archive/net/NNStat_3.3beta.tar.Z`.

Patents

A mechanism for arbitrating client access to a networked print server; US Patent 5,014,221 issued 7 May 1991

An Efficient Method Supporting Access to I/O Devices Through Large Direct-Mapped Data Caches; filed January 1990; US Patent 5,253,353 issued 12 October 1993.

Cache memory system and method for selectively removing stale aliased entries; filed June 1992; US Patent 5,675,763 issued 7 October 1997.

System for sampling network packets by only storing the network packet that its error check code matches with the reference error check code; filed May 1995; US Patent 5,790,799 issued 4 August 1998.

Method for Predictive Prefetching of Information Over a Communications Network; filed May 1995; US Patent 5,802,292 issued 1 September 1998.

Compression Protocol with Multiple Preset Dictionaries (with Andrei Broder and Michael Mitzenmacher); filed October 29, 1997; US Patent 5,953,503 issued 14 September 1999.

Computer Data Restoration Assembly and Associated Method; filed December 19, 1997; US Patent 6,052,764 issued 18 April 2000.

Method and apparatus of improving network performance and network availability in a client-server network by transparently replicating a network service; US Patent 6,097,882 issued 1 August 2000.

Method for Dynamically Adapting Web Content and Presentation to Network Conditions (with Lawrence Brakmo); US Patent 6,243,761 issued 5 June 2001.

System and Method for Reducing Latencies while Translating Internet Host Name-Address Bindings; US Patent 6,262,987 issued 17 July 2001.

System and method for receiver-based allocation of network bandwidth; US Patent 6,560,243 issued 6 May 2003.

System and Method for Allocating Server Output Bandwidth; US Patent 6,647,419 issued 11 November 2003.

Explicit server control of transcoding representation conversion at a proxy or client location; filed February 8, 2000. US Patent 6,704,798 issued 9 March 2004, US Patent 7,287,094 issued 23 October 2007

Analysis of Causal Relations Between Intercommunicating Nodes (with Marcos K. Aguilera); filed June 23, 2003. US Patent 7,254,646 issued 7 August 2007.

Reduction of network retrieval latency using cache and digest; filed April 3, 2001. US Patent 7,315,884 issued January 1, 2008.

Method and system for limiting the impact of undesirable behaviour of computers on a shared data network (with Michael Burrows and Raymond P. Stata); filed November, 2001. US Patent 7,383,574 issued June 3, 2008.

System and method for energy efficient data prefetching (with Keith Farkas, Parthasarathy Ranganathan, and Eduardo Pinheiro); filed December, 2001. US Patent 7,437,438 issued October 14, 2008.

Truncating Data Units; filed December 1, 2004; US Patent 7,512,705 issued March 31, 2009.

Method For Efficiently Detecting Node Addresses; filed 29 Oct 2008; US Patent 7,995,595 issued August 9, 2011.

Fast Detection Of Path Failure For TCP (With Mallikarjun B Chadalapaka, Randy Haagens, Richard Wrenn, Rick Wagner); filed July 28, 2006; US Patent 8,031,617 issued October 4, 2011

Main Memory With Non-volatile Memory and DRAM; Jeffrey Mogul, Eduardo Argollo de Oliveira Dias Junior, Paolo Faraboschi, and Mehul A. Shah; filed ca. 18 Sept 2009; US Patent 8,296,496 issued October 23, 2012

Proactive Sending Of An IP-to-MAC Address Binding For A High Fan-in Node; Jeffrey Mogul, Praveen Yalagandula, and Jayaram Mudigonda; filed 20 April 2010; US Patent 8,359,376 issued January 22, 2013.

System and Method for Ascribing Resource Consumption to Activity in a Causal Path of a Node of a Distributed Computing System; Jeffrey Mogul, Janet L. Wiener, Marcos K. Aguilera, Keith I. Farkas, and Parthasarathy Ranganathan; filed September 24, 2004; US Patent 8,364,829 issued January 29, 2013.

Distributing decision making in a centralized flow routing system; Jean Tourrilhes, Praveen Yalagandula, Puneet Sharma, Jeffrey C. Mogul, and Sujata Banerjee; filed May 10, 2010; US Patent 8,503,307 issued August 6, 2013.

Switching to core executing OS like codes upon system call reading greater than predetermined amount of data; Nathan L. Binkert, Jeffrey C. Mogul, Jayaram Mudigonda, and Parthasarathy Ranganathan; filed April 25, 2007; US Patent 8,886,917 issued November 11, 2014.

Methods and apparatus to access memory using runtime characteristics; Jeffrey Clifford Mogul, Naveen Muralimanohar, Mehul A. Shah, Eric A. Anderson; filed October 31, 2011; US Patent 9,146,867 issued September 29, 2015.

Generating network topologies; Jayaram Mudigonda, Praveen Yalagandula, Jeffrey Clifford Mogul; filed October 31, 2011; US Patent 9,148,348 issued September 29, 2015.

Replaceable modular optical connection assembly; Jeffrey Clifford Mogul, Dwight L. Barron, Kevin B. Leigh, Jayaram Mudigonda, Praveen Yalagandula, Guodong Zhang, Terrel Morris; filed March 12, 2012; US Patent 9,164,250 issued October 20, 2015.

Inferring causal paths in a distributed computing environment; Patrick A. Reynolds, Janet L. Wiener, Marcos K. Aguilera, Jeffrey C. Mogul; filed June 9, 2005; US Patent 9,178,721 issues November 3, 2015.

Forwarding broadcast traffic to a host environment; Jeffrey Mogul, Praveen Yalagandula, Jayaram Mudigonda; filed June 24, 2010; US Patent 9,191,328 issued November 17, 2015

Scalable multi-tenant network architecture for virtualized datacenters; Jayaram Mudigonda,

Praveen Yalagandula, Jeffrey C. Mogul, Bryan Stiekes, Anna Fischer; filed 7 June 2011; US Patent 9,304,798 issued April 5, 2016

Implementing a software defined network using event records that are transmitted from a network switch; Jeffrey C. Mogul, Paul T. Congdon, Dwight L. Barron; filed 29 June 2012; US Patent 9.325,569 issued April 26, 2016

Generating EGFT network topologies; Jayaram Mudigonda, Praveen Yalagandula, Jeffrey Clifford Mogul; filed 31 October 2011; US Patent 9,385,922 issued July 5, 2016

Reducing TCP timeouts due to Incast collapse at a network switch; Jeffrey Clifford Mogul, Paul T. Congdon, Jayaram Mudigonda, Lucian Popa filed 31 October 2011; US Patent 9,590,909 issued March 7, 2017

Method and system for enforcing multiple rate limits with limited on-chip buffering; Jeffrey Clifford Mogul, Jakov Seizovic, Benjamin Charles Seribrin; file 26 February 2015; US Patent 9,755,978 issued 5 September 2017

Method and system for validating rate-limiter determination made by untrusted software; Jeffrey Clifford Mogul, Jakov Seizovic, Yuhong Mao, Benjamin Charles Seribrin; file 26 February 2015; US Patent 9,762,502 issued 12 September 2017

Methods and apparatus to control generation of memory access requests; Jeffrey Clifford Mogul, Jayaram Mudigonda filed 31 October 2011; US Patent 9,772,958 issued 26 September 2017

Bandwidth guarantee and work conservation; Lucian Popa, Praveen Yalagandula, Sujata Banerjee, Jeffrey C. Mogul, Yoshio Turner, Jose Renato G. Santos; filed 31 July 2012; US Patent 9,794,185 issued 17 October 2017

Assigning selected groups to routing structures; Praveen Yalagandula, Jayaram Mudigonda, Jeffrey Clifford Mogul; filed 12 April 2012; US Patent 9,813,328 issued 7 November 2017

Resource Allocator; Jeffrey Clifford Mogul, Alvin AuYoung, Sujata Banerjee, Jung Gun Lee, Jean Tourrilhes, Michael Schlansker, Puneet Sharma, Lucian Popa; filed 20 July 2013; US Patent 10,009,285 issued 26 June 2018

Implementing a software defined network based on event records; Jeffrey C. Mogul; Paul T. Congdon; Dwight L. Barron; filed 29 June 2012; US Patent 10,257,084 issued 9 April 2019

Network multi-level rate limiter; Yuhong Mao; Richard Lee Sites; Jeffrey Clifford Mogul; filed 26 Feb 2015; US Patent 10,469,404 issued 5 November 2019

Governing bare metal guests; Jeffrey Clifford Mogul; Jose Renato G. Santos; Yoshio Turner; Kevin T. Lim; filed 30 Apr 2013; US Patent 10,728,171 issued 28 July 2020

Another 20+ patent applications have been filed but not yet granted or denied.

Keynotes and Invited Talks

"Performance Implications of Multiple Pointer Sizes," presented at the 1995 German UNIX User's Group Conference.

"What's wrong with HTTP, and why it doesn't matter," presented at the 1999 USENIX Technical Conference.

"Faster servers, slower Web?," keynote for the Performance Aspects of Web Servers workshop, in conjunction with SIGMETRICS 2001.

"Reducing energy consumption for networked applications," keynote for ACM/IEEE Symposium on Architectures for Networking and Communications Systems 2007. Also given at Juniper Networks Distinguished Speaker Series (Sep. 2008) and Intel Research Berkeley Lab (Nov. 2008).

"SIGOPS to SIGARCH: 'Now it's our turn to push you around' " Research Vision Session, OSDI 2010. Longer version given at UC Berkeley (Nov. 2010).

"Research Challenges for Modern Data-Center Networks," keynote for 23rd International Teletraffic Congress 2011. Also given at Technicolor Research Palo Alto (Sep. 2011)

"Corybantic: Towards the Modular Composition of SDN Controllers" given at Max Plank Institute for Software Systems (Nov. 2012); EPFL (Nov. 2012); University of Southern California (Feb. 2013).

Graduate Student Committees

Emre Kiciman, Stanford University Computer Science Department (Ph.D. granted 2005)

David Nellans, University of Utah (Ph.D. expected 2010)

Katrina LaCurts, Massachusetts Institute of Technology (Ph.D. expected 2014)

Professional Service

Journal Editor:

- *Internetworking: Research and Experience*, Associate Editor, 1994-1996

Program Committee Chair:

- USENIX Winter '94 Technical Conference
- IEEE TCOS Sixth Workshop on Hot Topics in Operating Systems
- Second Workshop on Industrial Experiences With Systems Software
- Workshop on Organizing Workshops, Conferences, and Symposia for Computer Systems (WOWCS 2008)
- ACM SIGCOMM Workshop: Research on Enterprise Networking (WREN 2009)

Program Committee Co-Chair:

- ACM SIGCOMM Workshop on Network-I/O Convergence: Experience, Lessons, Implications (NICELI 2003)
- USENIX/ACM Seventh Symposium on Operating Systems Design and Implementation (OSDI 2006)
- 6th ACM/IEEE Symposium on Architectures for Networking and Communications Systems (ANCS 2010)
- ACM SIGCOMM 2011
- Tenth USENIX Symposium on Networked Systems Design and Integration (NSDI 2013)

Conference Steering Committees:

- OSDI Steering Committee (2008-2010)
- ANCS Steering Committee (2010)
- HotICE Workshop Steering Committee (2011–2012)
- SIGCOMM conference Technical Steering Committee (2012–2015; chair, 2012-2013)
- SIGCOMM HotNets workshop Steering Committee (2013–2015; chair, 2014-2015)
- NSDI Steering Committee (2016–)

Award Committees:

- Member, ACM SIGOPS Mark Weiser Award Committee (2002-2004; Chair 2004)

- Member, ACM SIGOPS "Hall of Fame" Award Committee (2006-2012; Chair 2007; 2017); "Catch-up" Committee (2015)
- ACM SIGCOMM SOSR Software Systems Award (2017)
- Member, IEEE Internet Award Committee (2017–)

Additional Conference Organization Duties:

- Publicity Chair: 3rd IEEE-TCOS Workshop on Workstation Operating Systems

Program Committee Member:

- ACM SIGCOMM '88 Conference
- ACM SIGCOMM '92 Conference
- USENIX Summer '93 Technical Conference
- 1994 USENIX Symposium on High-Speed Networking
- ACM SIGCOMM '95 Conference
- ACM SIGCOMM '96 Conference
- USENIX/ACM Second Symposium on Operating Systems Design and Implementation
- 16th ACM Symposium on Operating Systems Principles
- Second USENIX Symposium on Internet Technologies and Systems
- 17th ACM Symposium on Operating Systems Principles
- IEEE Symposium on Applications and the Internet
- Sixth International Workshop on Web Caching and Content Distribution
- SIGCOMM Internet Measurement Workshop 2001
- Twelfth International World Wide Web Conference
- 9th Workshop on Hot Topics in Operating Systems
- USENIX/ACM Sixth Symposium on Operating Systems Design and Implementation
- 2nd Symposium on Networked Systems Design and Implementation (NSDI 2005)
- ACM SIGCOMM 2005 Conference
- 21st ACM Symposium on Operating Systems Principles (SOSP 2007)
- 2008 USENIX Annual Technical Conference
- 6th Symposium on Networked Systems Design and Implementation (NSDI 2009)
- 22nd ACM Symposium on Operating Systems Principles (SOSP 2009)
- ACM SIGCOMM 2010 Conference
- ACM SIGCOM Workshop on Programmable Routers for Extensible Services of Tomorrow (PRESTO 2010)
- Poster and Demo Program Committee, 23rd ACM Symposium on Operating Systems Principles (SOSP 2011)
- EuroSys 2011

- ACM SIGCOMM Workshop on Hot Topics in Software-Defined Networking (HotSDN 2012)
- 10th USENIX Symposium on Operating Systems Design and Implementation (OSDI 2012)
- Open Network Summit (ONS) 2013 Research Track
- ACM SIGCOMM Workshop on Hot Topics in Software-Defined Networking (HotSDN 2014)
- 11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 2014, External Review Committee)
- 25nd ACM Symposium on Operating Systems Principles (SOSP 2015)
- ACM SIGCOMM 2016 Conference
- Sixteenth ACM Workshop on Hot Topics in Networks (HotNets 2017)
- Eighteenth ACM Workshop on Hot Topics in Networks (HotNets 2019)

Journal Reviews:

- ACM Transactions on Computer Systems
- ACM SIGCOMM Computer Communication Review
- Communications of the ACM
- IEEE Communications Magazine
- IEEE Computer Magazine
- IEEE Journal on Selected Areas in Communications
- IEEE Network Magazine
- IEEE Transactions on Computers
- IEEE/ACM Transactions on Networking
- IEEE Transactions on Software Engineering
- Journal of Internetworking
- Software Practice and Experience
- Digital Technical Journal

Additional Conference Reviews:

- Fourth Conference on Architectural Support for Programming Languages and Operating Systems
- Sixth Conference on Architectural Support for Programming Languages and Operating Systems
- Seventh Conference on Architectural Support for Programming Languages and Operating Systems
- 20th International Symposium on Computer Architecture
- 23rd International Symposium on Computer Architecture
- Twelfth Symposium on Operating Systems Principles
- USENIX Winter '93 Technical Conference

- USENIX/SIGOPS 1994 Operating Systems Design and Implementation Conference
- USENIX Winter '95 Technical Conference
- SIGMETRICS 1998, 1999
- SIGCOMM 1998

Government Committees:

- National Science Foundation, pre-proposal reviewer, 2000
- National Science Foundation Computer & Information Science & Engineering (CISE) Infrastructure Advisory Subcommittee, 2012-2013

Professional Society Memberships: Association for Computing Machinery; ACM SIGOPS; ACM SIGCOMM; IEEE; Sigma Xi.

Standards Committee Participation

- Chair, Internet Engineering Task Force (IETF) MTU Discovery Working Group
- IETF Router Discovery Working Group
- IETF HTTP Working Group
- Chair, informal working group on Pulse Per Second API

Technical Skills

Proficient
C, Awk
kernel and user-level programming: BSD-based UNIX, Linux

Formerly Proficient
Pascal, Modula-2, FORTRAN, assembly code
kernel and user-level programming: VAX/VMS, RT-11

Familiar
Perl, PHP

Miscellaneous

Citizenship: United States of America

Erdös Number: ≤ 4

References

Available upon request

September 8, 2020