

James Hendricks

Computer Science Department, Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA 15213
(412) 849-8145 • James.Hendricks@cs.cmu.edu

Objective

To research and develop new technologies that ensure consumers can efficiently access whatever data and applications they need, wherever they may be, and whenever they wish, despite asynchronous networks, weak connectivity, and the inevitable faults and corruptions caused by modern hardware and software.

Education

Ph.D. in Computer Science, Carnegie Mellon University *Spring 2009 (expected)*

- Dissertation Topic: Efficient Byzantine Fault Tolerant Storage and Services
- Advisor: Professor Greg Ganger

M.S. in Computer Science, Carnegie Mellon University *Spring 2006*

B.S. in Electrical Engineering and Computer Science, UC Berkeley *May 2002*

- Honors focus in Mathematics

Research Interests

I have designed and built several distributed systems over the past few years. Most recently, I built an erasure-coded distributed storage system and a fault-tolerant replicated state machine. In general, I am interested in the applied and experimental aspects of computer science, including distributed systems, security, storage, and networking, and in the design and deployment of real systems.

Brief Employment History

Graduate Research Assistant, Carnegie Mellon University	<i>9/2002–Present</i>
Intern, IBM Almaden Storage Systems Group	<i>6/2006–8/2006</i>
Intern, Cisco Gigabit Systems Business Unit	<i>6/2002–8/2002</i>
Intern, Los Alamos National Laboratory	<i>Summers 1998–2001</i>

Awards and Honors

- Best Paper Award, 4th USENIX Conference on File and Storage Technologies (FAST 2005)
- U.S. National Defense Science and Engineering Graduate Fellowship (NDSEG), 2003-2006
- U.S. National Science Foundation (NSF) Graduate Fellowship, 2002-2003, 2006-2007
- Tau Beta Pi (officer) and Phi Beta Kappa
- Goldwater Scholar in Mathematics, Science, and Engineering, 2000-2002

Employment History

Carnegie Mellon University. *Dissertation research.*

Fall 2006–Present

For my dissertation, I am developing techniques for building efficient file systems and other distributed services that tolerate arbitrary (Byzantine) faults.

Spring 2008–Present: Developed Zzyzx, a protocol for scalable distributed services that tolerate arbitrary (Byzantine) faults. Zzyzx uses Byzantine Locking to perform substantially better than prior protocols. Zzyzx offers true scalability, in that adding servers increases throughput proportionally.

Spring 2008–Present: Defined correctness for protocols that tolerate faulty clients and proved when different notions of correctness are possible.

Fall 2006–Fall 2007: Built an erasure-coded Byzantine fault-tolerant block storage prototype that performs substantially better than prior approaches due to novel techniques, including homomorphic fingerprinting and partial encoding. (See my SOSP 2007 paper.)

Fall 2006–Summer 2007: Invented homomorphic fingerprinting, a cryptographic primitive that extends hashing to erasure-coded data. (See my PODC 2007 paper.)

Carnegie Mellon University. *Non-dissertation research.*

9/2002–Present

Fall 2008–Present: Researched lightweight authentication in protocols such as AMQP.

Spring 2005–Spring 2006: Researched scalable metadata for cluster-based storage.

Fall 2005–Fall 2006: Researched methods for cluster IO benchmarking. (See my FAST 2007 paper.)

Spring 2004–Fall 2005: Developed the prototype Self-* Storage system. (See my FAST 2005 paper.)

Spring 2003–Fall 2004: DNS Measurement. (See my IMC 2004 paper.)

Fall 2002–Fall 2004: Researched trusted computing and secure architectures. (See my SIGOPS '04 paper.)

IBM Almaden. *Intern, Storage Systems Group.*

6/2006–8/2006

Developed a distributed storage system research prototype.

Cisco Systems. *Intern, Gigabit Systems Business Unit.*

6/2002–8/2002

Fixed bugs and developed features for the Catalyst 4000 series switch. Learned about software architecture (IOS, CatOS) and hardware architecture of Cisco switches.

Los Alamos National Laboratory. *Intern, Advanced Computing Lab.*

Summers 1998–2001

Summer 2001: Verilog prototype of a high performance network interface.

Summer 2000-2001: LinuxBIOS SMP support, video and RAM initialization, flash write sequence.

Winter 2000: Built a proof-of-concept of LinuxBIOS (LinuxBIOS is now coreboot).

Summer 1999: Built a cluster benchmarking tool.

Teaching Experience

Teaching Assistant, Carnegie Mellon University

8/2007–12/2007

Advanced Operating Systems and Distributed Systems (15-712), for Professor David Andersen.

Teaching Assistant, Carnegie Mellon University

8/2003–12/2003

Advanced Operating Systems and Distributed Systems (15-712), for Professor Greg Ganger.

Publications

Refereed Publications:

1. James Hendricks, Gregory R. Ganger, Michael K. Reiter. Low-overhead Byzantine fault-tolerant storage. In *Proceedings of the Twenty-First ACM Symposium on Operating Systems Principles (SOSP 2007)*, Stevenson, WA, October 2007.
2. James Hendricks, Gregory R. Ganger, Michael K. Reiter. Verifying distributed erasure-coded data. In *Proceedings of the Twenty-Sixth Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC 2007)*, Portland, August 2007.
3. Michael Mesnier, Matthew Wachs, Raja R. Sambasivan, Julio Lopez, James Hendricks, Gregory R. Ganger, David O'Hallaron. //TRACE: Parallel trace replay with approximate causal events. In *Proceedings of the 5th USENIX Conference on File and Storage Technologies (FAST 2007)*, San Jose, February 2007.
4. Michael Abd-El-Malek, William V. Courtright II, Chuck Cranor, Gregory R. Ganger, James Hendricks, Andrew J. Klosterman, Michael Mesnier, Manish Prasad, Brandon Salmon, Raja R. Sambasivan, Shafeeq Sinnamohideen, John D. Strunk, Eno Thereska, Matthew Wachs, Jay J. Wylie. Early experiences on the journey towards self-* storage. *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering*, September 2006.
5. Michael Abd-El-Malek, William V. Courtright II, Chuck Cranor, Gregory R. Ganger, James Hendricks, Andrew J. Klosterman, Michael Mesnier, Manish Prasad, Brandon Salmon, Raja R. Sambasivan, Shafeeq Sinnamohideen, John D. Strunk, Eno Thereska, Matthew Wachs, and Jay J. Wylie. Ursa Minor: Versatile Cluster-based Storage. In *Proceedings of the 4th USENIX Conference on File and Storage Technologies (FAST 2005)*, San Francisco, December 2005. (Best Paper Award)
6. Jeffrey Pang, James Hendricks, Aditya Akella, Bruce Maggs, Roberto De Prisco, and Srinivasan Seshan. Availability, Usage, and Deployment Characteristics of the Domain Name System. In *Proceedings of the Second ACM/USENIX Internet Measurement Conference (IMC 2004)*, Taormina, Sicily, October 2004.
7. James Hendricks and Leendert van Doorn. Secure Bootstrap is Not Enough: Shoring up the Trusted Computing Base. In *Proceedings of the Eleventh SIGOPS European Workshop (SIGOPS 2004)*, ACM SIGOPS, Leuven, Belgium, September 2004.
8. Ron Minnich, James Hendricks, and Dale Webster. The Linux BIOS. In *Proceedings of The Fourth Annual Linux Showcase and Conference (ALS 2000)*, Atlanta, GA, October 2000. USENIX Association.

Unpublished (in preparation):

9. James Hendricks, Gregory R. Ganger, Michael K. Reiter. Zzyzx: Scalable Fault Tolerance Through Byzantine Locking. In preparation. Preliminary work accepted for presentation at the Work-in-Progress Session and the Poster Session at OSDI 2008.
10. James Hendricks, Gregory R. Ganger, Michael K. Reiter. The Correctness of Distributed Systems in the Presence of Faulty Clients. In preparation.

Technical Reports:

11. Shafeeq Sinnamohideen, James Hendricks, Raja R. Sambasivan, and Gregory R. Ganger. Eliminating cross-server operations in scalable file systems. Carnegie Mellon University Parallel Data Lab Technical Report CMU-PDL-06-105, May, 2006.
12. James Hendricks, Raja R. Sambasivan, Shafeeq Sinnamohideen, and Gregory R. Ganger. Improving small file performance in object-based storage. Carnegie Mellon University Parallel Data Lab Technical Report CMU-PDL-06-104, May, 2006.

Professional Service and Affiliations

External reviewer for several conferences, including DSN, FAST, OSDI, Supercomputing, USENIX Technical, ASPLOS, Mobisys, and USENIX Security. Member of IEEE Computer Society, ACM SIGOPS, USENIX.

References

Professor Gregory Ganger
Director, Parallel Data Lab
Carnegie Mellon University

Dr. Leendert van Doorn
Adjunct Professor, Rice University
Senior Director and Senior Fellow, AMD

Professor Michael Reiter
Lawrence M. Slifkin Distinguished Professor
University of North Carolina at Chapel Hill

Professor Adrian Perrig
Technical Director, CyLab
Carnegie Mellon University