

# Andrey Ermolinskiy

## Contact Information:

*andrey@cs.berkeley.edu*

<http://www.cs.berkeley.edu/~andrey>

Tel: (510)-325-0045

## School Address:

464 Soda Hall, U.C. Berkeley,  
Berkeley, CA 94720

## EDUCATION

*University of California - Berkeley (Berkeley, CA; 2005 – present)*

- Ph.D. in Computer Science (*EECS Department*)
  - *Major field of study:* Distributed Systems and Networking.
  - *Minor field of study:* Database Management Systems.
- Certificate in Management of Technology (*Haas School of Business; June 2009*).
- M.S. in Computer Science (*EECS Department; June 2007*).

*Princeton University (Princeton, NJ; 1998 - 2002)*

- B.S.E. in Computer Science. High Honors academic recognition.
- Elected to the Society of Sigma XI for excellence in research.

*Akranes High School (Akranes, Iceland; 1996 - 1998)*

- Ranked first in the graduating class.
- Awarded the Akranes city scholarship for excellence in the natural sciences.
- Ranked 5<sup>th</sup> in the National Physics Olympiad, 1998.

## RESEARCH FOCUS

- Broadly, my research interests lie in the areas of hypervisor-based security, enterprise data management, high-performance distributed storage systems, and Internet architecture.
- My Ph.D. thesis investigates Practical Data Confinement [1] – a novel hypervisor-driven information security architecture for enterprise environments. PDC enables organizations to track the movement of confidential information, enforce dissemination restrictions, and protect sensitive documents and datasets against theft. PDC relies on fine-grained (instruction-level) information flow tracking techniques and achieves efficiency through speculative execution and parallelization. Unlike prior efforts in this area, PDC requires no changes to applications or the operating system, and is thus readily deployable in existing IT environments.
- In [2], I investigated the problem of concurrent access coordination for shared-disk parallel applications. I conceived and led the development of a novel synchronization primitive that lifts the safety and liveness limitations associated with the traditional approaches based on conservative distributed locking.
- My earlier work focused on developing and evaluating robust protocols for inter-domain routing [5, 6, 8] and exploring alternative Internet architectures [7].

## EMPLOYMENT HISTORY

*IBM Almaden Research Center (storage group) – Research Intern (San Jose, CA; 06/2007 – 09/2007).*

- Research on algorithms and techniques for lazy replica synchronization and cooperative caching in widely-distributed GPFS/pNFS environments.

*Intel Research – Research Intern (Berkeley, CA; 06/2006 – 09/2006)*

- Research on algorithms and tools for use in forensic analysis of distributed denial-of-service (DDoS) attacks against network servers.

*IBM Corporation – Software Engineer/Architect (Poughkeepsie, NY; 09/2002 – 08/2005)*

### **High-performance I/O division - General Parallel File System (GPFS)**

- Architected and implemented support for disaster-resilient GPFS clusters to address availability requirements for several key customers. The design is based on logged synchronous mirroring of data and metadata across a pair of geographically separated sites, using a third-site node as a tiebreaker for consensus protocols.
- Designed and implemented a number of other features and improvements in the areas of I/O performance, distributed token management, transaction recovery, and cluster configuration management.

**Sun Microsystems** – Intern Software Engineer (*Menlo Park, CA; 06/2001 – 09/2001*)

**Solaris kernel development**

Designed and implemented several kernel extensions for the Solaris platform including:

- A resident set monitor for memory-intensive workloads.
- A page pre-faulting mechanism to reduce start-up times of large software applications.

**IBM Corporation** – Intern Software Engineer (*Austin, TX; 06/2000 – 09/2000*)

**High-performance I/O division - General Parallel File System (GPFS)**

- Designed and implemented an optimized kernel memory allocation mechanism for GPFS on Linux.
- Implemented error logging support for GPFS on Linux.
- Helped on porting other parts of the project from AIX to Linux and configuring an RS/6000 ScalablePOWERparallel (SP) system.

**HB International** - Intern Software Engineer (*Reykjavik, Iceland; 06/1999 - 09/1999*)

**Centara development**

- Designed and implemented a protocol for EFT (Electronic Fund Transfer) on top of NetBIOS.

**High Speed Information Inc.** – Intern Software Engineer (*Reykjavik, Iceland; 06/1998-09/1998*)

- Implemented and tested parts of an Aeronautical Decision Support System using Active Server Pages.
- Implemented an ActiveX-based automated database access solution for the Technical Support department.

## SKILLS SUMMARY

**Programming Languages:** C/C++, x86 assembly, Java, ML, MS Visual Basic, HTML, SQL.

**User-level development:** Linux, AIX, Solaris, and Windows.

**Kernel-level development:** Linux and AIX; Significant experience with the internals of the Linux kernel (in particular its storage and networking stacks).

**Hypervisor and emulator development:** Strong familiarity with the internals of Xen and QEMU. Design and implementation of mechanisms for instruction-level information flow tracking, dynamic code generation, and security policy enforcement.

**Concepts, algorithms, and protocols:** Strong familiarity with the theoretical and practical aspects of developing robust, scalable, and fault tolerant distributed systems; Expert understanding of distributed protocols (consensus, leader election, mutual exclusion, state machine replication, atomic commit, replica synchronization, transaction recovery, etc); Expert understanding of parallel and distributed filesystem architecture. Expert understanding of Internet transport and routing protocols (TCP/IP, BGP, OSPF, RIP, DVMRP, PIM-SM, PIM-DM); Expert understanding of file- and block-level remote storage protocols (NFS, iSCSi).

**Other:** Fluent in Icelandic, English, Russian.

## SELECTED PUBLICATIONS

- [1] “*Practical Data Confinement*”. A. Ermolinskiy, S. Katti, S. Shenker, L. Fowler, M. McCauley. (In submission, Oct 2009)
- [2] “*Minuet: Rethinking Concurrency Control in Storage Area Networks*”. A. Ermolinskiy, D. Moon, B. G. Chun, S. Shenker. *Proceedings of FAST’09* (Feb. 2009).
- [3] “*c2cfs – a Collective Caching Architecture for Distributed File Access*”. A. Ermolinskiy, R. Tewari. *Proceedings of NSDM’09* (June 2009).
- [4] “*S3 – Securing Sensitive Stuff*”. S. Katti, A. Ermolinskiy, M. Casado, S. Shenker, H. Balakrishnan. *OSDI’08 Work-in-Progress (WiP) report* (December 2008).
- [5] “*Reducing Transient Disconnectivity using Anomaly-Cognizant Forwarding*”. A. Ermolinskiy, S. Shenker. *Proceedings of ACM HotNets-VII* (Oct. 2008).
- [6] “*The Design and Implementation of Free Riding Multicast*”. A. Ermolinskiy. *Master’s Report*. (May 2007).
- [7] “*A Data-Oriented (and Beyond), Network Architecture*”. T. Kopenon, M. Chawla, B. G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, I. Stoica. *Proceedings of ACM SIGCOMM’07* (Aug. 2007).
- [8] “*Revisiting IP Multicast*” S. Ratnasamy, A. Ermolinskiy, S. Shenker. *Proceedings of ACM SIGCOMM’06* (Sept. 2006).
- [9] “*Disaster Recovery with General Parallel File System*” IBM technical whitepaper (Aug. 2004).
- [10] “*Pitch Histograms in Audio and Symbolic Music Information Retrieval*” G. Tzanetakis, A. Ermolinskiy, P. Cook. *Proceedings of ISMIR 2002* (Oct. 2002).
- [11] “*Beyond the Query-by-Example Paradigm: New Query Interfaces for Music Information Retrieval*”. G. Tzanetakis, A. Ermolinskiy, P. Cook. *Proceedings of ICMC’02* (Sept. 2002).

## PROFESSIONAL SERVICE

- Outside reviewer for the following conferences: SRDS’08, PDCAT’09