Andrey Ermolinskiy, Ph.D.

E-mail: andrey@crimsonvoyage.com Web: http://www.crimsonvoyage.com

EDUCATION

University of California - Berkeley *(09/2005 – 06/2011)*

- Ph.D. in Computer Science (*EECS Department; 06/2011)*
  - *Major field of study:* Distributed Systems, Networking, Information Security
  - *Minor field of study:* Database Management Systems
- Certificate in Management of Technology (*Haas School of Business*; *05/2009*)
- M.S. in Computer Science (*EECS Department*; *06/2007)*, GPA: 3.9

Princeton University *(09/1998 – 05/2002)*

- B.S.E. in Computer Science with High Honors, GPA: 3.7

Akranes High School *(Iceland; 09/1996 – 06/1998)*

- Ranked first in the graduating class
- Awarded the Akranes city scholarship for excellence in the natural sciences

AWARDS AND HONORS

- Internationally-recognized expert in computer networking, distributed systems, operating systems, and virtualization with 8 research publications in top-tier academic conferences
- Winner of the 1st prize in the 2011 QuantCup programming competition (http://www.quantcup.org)
  - Competition participants were asked to design and implement a highly-optimized Price-Time Order Matching Engine (OME) in C++ for a simulated electronic stock exchange. Submissions were judged based on performance (order execution latency and jitter), correctness, and coding style. I discovered a new algorithmic technique for speeding up the order matching computation and implemented a highly-efficient OME on the basis of this technique. My implementation was awarded the 1st prize, having been recognized for its extreme efficiency.
- Elected to the Society of Sigma Xi for excellence in undergraduate research *(Princeton University, 05/2002)*
- 5th place in the National Physics Olympiad *(Reykjavik, Iceland, 08/1998)*

RESEARCH FOCUS

My research expertise spans a number of areas in the general field of computer systems, focusing specifically on distributed systems, networking, storage systems, operating systems, virtual machines, hypervisors, optimizing compilers, and information flow analysis. My Ph.D. thesis [9] investigates Practical Information Flow Tracking (PIFT) - a novel hypervisor-driven information security platform for enterprise environments. In broad terms, PIFT provides mechanisms for tracking the movement of sensitive information between applications, machines, and users, allowing enterprise security administrators to enforce high-level information flow policies and protect sensitive documents from unauthorized disclosure to external parties. PIFT relies on fine-grained (instruction-level) information flow tracking techniques and achieves efficiency through speculative execution and parallelization. Unlike numerous prior research efforts in this area, PIFT requires no changes to applications or the operating system.

In the past, I have worked on developing novel synchronization protocols for distributed storage systems [1,2], investigating new interdomain packet routing algorithms [4,6], and exploring new architectures for future Internet [5].

EMPLOYMENT HISTORY

Headlands Technologies *07/2013 - Present*

***Qualcomm Research Silicon Valley*:** Staff Research Engineer *03/2012 – 07/2013*
Senior Research Engineer *03/2011 - 03/2012*

- Conducting research on virtual machines, runtime compilers, and hypervisors for future mobile computing devices. Currently leading a team of 4 researchers and engineers in designing and implementing an advanced Java virtual machine to allow unmodified Android applications to exploit the full capabilities of Qualcomm's heterogeneous processor architecture. Our project has designed, implemented, and transferred state-of-the-art VM and runtime code generation technology to product groups in order to enhance the performance of Qualcomm devices.

***Qualcomm Research Silicon Valley*:** Research Intern *06/2010 – 08/2010*

- Design and implementation of parallel scheduling algorithms and managed runtime environments for future mobile devices.

IBM Research – Almaden (Storage systems group): Research Intern *06/2007 – 09/2007*

- Research on algorithms and techniques for lazy replica synchronization and cooperative caching in geographically-distributed NFS/GPFS configurations.

Intel Research – Berkeley: Research Intern *06/2006 – 09/2006*

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

***IBM Corporation*:** Software Engineer *09/2002 – 08/2005*

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

- Designed and implemented support for disaster recovery in GPFS clusters to address high-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: Software Engineer (Intern) *06/2001 – 09/2001*

Solaris kernel group

- Designed and implemented several kernel extensions for the Solaris operating system, including: (1) A resident set monitor for memory-intensive workloads, (2) A page pre-faulting mechanism to reduce start-up times of large software applications.

***IBM Corporation*:** Software Engineer (Intern) *06/2000 – 09/2000*

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

- Designed and implemented an optimized kernel memory allocation mechanism for GPFS on Linux.

TECHNICAL EXPERTISE

- **Design and implementation of distributed systems**: Strong familiarity with the theoretical and practical aspects of developing robust, scalable, and fault tolerant distributed systems. Expert understanding of fundamental distributed computing primitives (consensus, state machine replication, reliable multicast, etc.).
- **Storage systems**: Expert understanding of filesystem architecture, including parallel and distributed filesystems.
- **Operating systems:** Significant experience with the internals of the Linux/Android kernel.
- **Virtual machines, hypervisors, emulators:** Significant experience with the internals of Xen and QEMU. Expert understanding of the architecture and implementation of the Dalvik VM.
- **Compiler development:** Experience with implementing compiler transformations (global value numbering, conditional constant propagation, copy coalescing) and graph-coloring register allocators.
- **Programming languages:** C/C++, x86 assembly, ARM assembly, Java.

PUBLICATION LIST

Conference Proceedings

[1] "C2Cfs: A Collective Caching Architecture for Distributed File Access". Andrey Ermolinskiy, Renu Tewari. In *Proceedings of the 2009 International Workshop on Network Storage and Data Management (NSDM'09)*, Seoul, Korea, pp 642-647, June 2009.

[2] "Minuet: Rethinking Concurrency Control in Storage Area Networks". Andrey Ermolinskiy, Daekyeong Moon, Byung-Gon Chun, Scott Shenker. In *Proceedings of the 7th USENIX Conference on File and Storage Technologies (FAST'09)*, San Francisco, CA, pp 311-324, February 2009.

[3] "S3: Securing Sensitive Stuff". Sachin Katti, Andrey Ermolinskiy, Martin Casado, Scott Shenker, Hari Balakrishnan. *USENIX OSDI 2008 Work in Progress Report*, San Diego, CA, December 2008.

[4] "Reducing Transient Disconnectivity using Anomaly-Cognizant Forwarding". Andrey Ermolinskiy, Scott Shenker. In *Proceedings of the 7th ACM Workshop on Hot Topics in Networks (HotNets-VII)*, Calgary, Canada, pp. 91-96, October 2008.

[5] "A Data-Oriented (and Beyond), Network Architecture". Teemu Koponen, Mohit Chawla, Byung-Gon Chun, Andrey Ermolinskiy, Kye Hyun Kim, Scott Shenker, Ion Stoica. In *Proceedings of ACM Special Interest Group on Data Communications Conference (SIGCOMM 2007)*, Kyoto, Japan, August 2007.

[6] "Revisiting IP Multicast". Sylvia Ratnasamy, Andrey Ermolinskiy, and Scott Shenker. In *Proceedings of ACM Special Interest Group on Data Communications Conference (SIGCOMM 2006)*, Pisa, Italy, pp. 15-26, September 2006.

[7] "Pitch Histograms in Audio and Symbolic Music Information Retrieval". George Tzanetakis, Andrey Ermolinskiy, and Perry Cook. In *Proceedings of the Third International Conference on Music Information Retrieval (ISMIR 2002)*, Paris, France, pp. 31-38, October 2002.

[8] "Beyond the Query-by-Example Paradigm: New Query Interfaces for Music Information Retrieval". George Tzanetakis, Andrey Ermolinskiy, Perry Cook. In *Proceedings of the International Computer Music Conference (ICMC 2002)*, Gothenburg, Sweden, pp. 177-183, September 2002.

Technical Reports

[9] Ph.D. Thesis: “Design and Implementation of a Hypervisor-Based Platform for Dynamic Information Flow Tracking in a Distributed Environment". Andrey Ermolinskiy. *UCB Technical Report (UCB/EECS-2011-50)*, May 2011.

[10] “Towards Practical Taint Tracking". Andrey Ermolinskiy, Sachin Katti, Scott Shenker, Lisa Fowler, Murphy McCauley. *UCB Technical Report (UCB/EECS-2010-92)*, June 2010.

[11] “Practical Data Confinement". Andrey Ermolinskiy, Sachin Katti, Scott Shenker, Lisa Fowler, Murphy McCauley. *UCB Technical Report*, September 2009.

[12] "C2Cfs: A Collective Caching Architecture for Distributed File Access". Andrey Ermolinskiy, Renu Tewari. *UCB Technical Report (UCB/EECS-2009-40)*, March 2009.

[13] “Reducing Transient Disconnectivity using Anomaly-Cognizant Forwarding". Andrey Ermolinskiy, Scott Shenker. *UCB Technical Report (UCB/EECS-2008-120)*, September 2008.

[14] "Minuet: Rethinking Concurrency Control in Storage Area Networks". Andrey Ermolinskiy, Daekyeong Moon, Byung-Gon Chun, Scott Shenker. *UCB Technical Report (UCB/EECS-2008-57)*, May 2008.

[15] Master's Thesis: "Design and Implementation of Free Riding Multicast". Andrey Ermolinskiy. *UCB Technical Report,* May 2007.

[16] "An Experiment in Deploying Next Generation Network Protocols". Hitesh Ballani, Andrey Ermolinskiy, Sylvia Ratnasamy, Paul Francis. *Cornell University Technical Report (TR2006-2028),* May 2006.

[17] "Disaster Recovery with General Parallel File System". Andrey Ermolinskiy. *IBM Technical Whitepaper*, August 2004.

PROFESSIONAL SERVICE

Served as a reviewer for following conferences:

- *The 27th IEEE International Symposium on Reliable Distributed Systems (SRDS 2008)*
- *The 10th International Conference on Parallel and Distributed Computing (PDCAT 2009)*