

BYUNG-GON CHUN

International Computer Science Institute (ICSI)
1947 Center Street, Suite 600
Berkeley, CA 94704

1-510-225-8990
bgchun@cs.berkeley.edu
<http://www.cs.berkeley.edu/~bgchun>

RESEARCH INTERESTS

Distributed systems, networking, fault-tolerance, security, complexity, troubleshooting

EDUCATION

University of California, Berkeley, Berkeley, CA
Doctor of Philosophy, Computer Science, 2007

- Dissertation topic: *Mechanisms to Tolerate Misbehavior in Replicated Systems*
- Major areas: Systems, Networking; Minor areas: Algorithms, Statistics
- Advisors: Prof. Scott Shenker and Prof. John Kubiawicz ; GPA: 4.0 / 4.0

Stanford University, Stanford, CA

Master of Science, Computer Science, 2002

- Advisor: Dr. Mary Baker ; GPA: 3.94 / 4.0

Seoul National University, Seoul, Korea

Master of Science, Electronic Engineering, 1996

Bachelor of Science, Electronic Engineering, 1994

RESEARCH & PROFESSIONAL EXPERIENCE

International Computer Science Institute, Berkeley, CA 2007 – Present
Postdoctoral Researcher, Networking Group, working with Prof. Scott Shenker, Dr. Petros Maniatis, and Dr. Sylvia Ratnasamy

Diverse replication for single-machine Byzantine-fault tolerance: Investigate exploiting cores in many-core systems to defend against software attacks. Explore different isolation and software diversity mechanisms for single-machine Byzantine-fault tolerance. Explore the design space of Byzantine-fault tolerance protocols enabled by this new single-machine execution environment.

Distributed systems troubleshooting: Develop a distributed troubleshooting system called D3 that executes distributed online/offline queries on locally produced execution tracing and resource utilization logs. Queries are declaratively specified and are turned into data flow graphs using P2. Apply D3 to debug performance/fault problems of Hadoop, an open source MapReduce implementation.

Complexity of networked systems: Proposed a metric that quantifies the complexity of networked system designs. The complexity metric is derived from the ensemble of distributed state dependencies. Applied the metric to networked systems such as Internet routing, classical distributed systems, and resource discovery, and showed insights in designing systems with respect to complexity.

University of California, Berkeley, Berkeley, CA 2002 – 2007
Research Assistant, Computer Science Division

Byzantine-fault tolerance of distributed systems: Proposed a new small trusted primitive, Attested Append-only Memory (A2M), that provides commitment and solves equivocation

problems where a faulty node can tell different protocol messages to different nodes at the same protocol step. Constructed A2M-enabled Byzantine-fault tolerant (BFT) replicated systems that improve fault bounds; A2M-PBFT-EA guarantees both liveness and safety (i.e., linearizability) with up to one-half faulty nodes and A2M-PBFT-E guarantees safety up to two-thirds faulty nodes. Also designed A2M-Storage, a protocol that guarantees linearizability in an untrusted server thanks to A2M. Small trusted primitives like A2M have broader implications on structuring trustworthy systems to achieve better fault-tolerance properties.

Proposed a tiered Byzantine fault model where components in different tiers have different fault characteristics and applied it to a long-term service. Designed and implemented TimeMachine, a long-term BFT key-value store that correctly preserves non-self-certifying data despite short-term correlated faults (i.e., $N-1$ faulty replicas out of N total replicas). TimeMachine achieves its preservation property by integrating proactive recovery and a trusted device that contains rate-limited append-only memory and batching state changes during proactive recovery.

Distributed storage: Designed and analyzed efficient replica maintenance for durability of distributed storage systems such as DHTs. Applied queuing theory to model durability of replica maintenance and examined efficient replica repair and placement strategies using a PlanetLab trace. Instrumented slow node problems on PlanetLab and proposed routing techniques to improve *get* latencies of OpenDHT, a public DHT service. Also developed a scalable, safe DHT key consistency algorithm for OpenDHT. This algorithm guarantees at most one node is responsible for a specific key range.

Distributed consensus in inter-domain routing: Analyzed a distributed precedence-value passing algorithm that solves policy-induced BGP instability problems; the algorithm runs only when BGP dispute wheels inducing oscillations exist and resolves them without losing ISP's policy autonomy. Developed a simulator and demonstrated convergence using experiments with diverse network topologies.

Anycast: Designed and implemented ChunkCast, an anycast service for distributing large chunked-files. Chunkcast exploits both peer download simultaneity and peer locality. Also developed a distributed locality-aware BitTorrent tracker built atop ChunkCast and deployed it on PlanetLab. Designed a Data-Oriented Network Architecture (DONA), which is a "clean-slate" design for naming and name resolution focusing on the shift in usage from host-centric to data-centric applications. DONA provides authenticity using self-certifying names and availability using name-based anycast.

Game-theoretic analysis of federated systems: Proposed new game-theoretic models for selfish caching where each node stores replicas based on its payoff. Analyzed price of anarchy with and without payments and proved that payments can induce Nash equilibria that are socially optimal. Also characterized selfishly constructed overlay networks where each node sets up links based on its cost to route and analyzed properties of graphs generated by games.

Intel Research, Berkeley, CA

Summer 2006

Research Intern, worked with Dr. Petros Maniatis.

Digital preservation: Examined digital preservation techniques that emulate unbounded long-term write-once memory. Explored time-varying fault models and design tradeoffs for distributed audit and repair protocols.

Stanford University, Stanford, CA

2001 – 2002

Research Assistant, Computer Science Department, worked with Dr. Mary Baker *Scheduling in mobile ad hoc networks*: Examined queuing behavior of packets in mobile ad hoc networks and analyzed the effects of new packet scheduling algorithms on throughput and end-to-end delay with varying mobility and traffic load.

Sundosoftware Inc., Seoul, Korea 1996 – 1999

Software Engineer

Spatial database system: Designed fast spatial query and network algorithms for spatial database systems. Developed middleware server and client for an intelligent transportation system that perform geographic mapping, shortest path routing with routing constraints, traffic monitoring, and topological map editing.

Seoul National University, Seoul, Korea 1994 – 1996

Research Assistant, Electronic Engineering Department

TCP over cellular networks: Proposed a base station packet filter that eliminates duplicate packets and a new TCP timeout scheme not to waste wireless bandwidth.

TEACHING EXPERIENCE

Teaching Assistant, Computer Networks (EECS122), UC Berkeley Fall 2004

Stella: Designed a peer-to-peer file sharing system class project and implemented the prototype of the system. Students built both Chord and a distributed index built on Chord.

HONORS

- Samsung Lee Kun Hee Foundation Scholarship (2003-2006)
- Ilju Academic Foundation Scholarship (1999-2002)
- Graduated *magna cum laude* from Seoul National University (1994)

PROFESSIONAL ACTIVITIES

- Reviewer: ACM SOSP '07, USENIX/ACM NSDI '07 and '08, USENIX FAST '05, IEEE INFOCOM '05, IEEE P2P '05, DISC '07, IEEE Trans. on Parallel and Distributed Systems, IEEE/ACM Trans. on Networking, IEEE JSAC
- Founder and Organizer: UC Berkeley Systems Lunch Meeting 2004-2005

REFEREED PUBLICATIONS

- **Byung-Gon Chun**, Petros Maniatis, and Scott Shenker. *Diverse Replication for Single-Machine Byzantine-Fault Tolerance*. To appear in 2008 USENIX Annual Technical Conference (USENIX '08).
- **Byung-Gon Chun**, Sylvia Ratnasamy, and Eddie Kohler. *A Complexity Metric for Networked System Designs*. To appear in Fifth USENIX Symposium on Networked Systems Design and Implementation (NSDI '08).
- **Byung-Gon Chun**, Petros Maniatis, John Kubiatowicz, and Scott Shenker. *Attested Append-Only Memory: Making Adversaries Stick to their Word*. Twenty First ACM Symposium on Operating Systems Principles (SOSP '07), October 2007.
- Teemu Koponen, Mohit Chawla, **Byung-Gon Chun**, Andrey Ermolinskiy, Kye Hyun Kim, Scott Shenker, and Ion Stoica. *A Data-Oriented (and Beyond) Network Architecture*. ACM SIGCOMM, August 2007.

- Cheng Tien Ee, **Byung-Gon Chun**, Vijay Ramachandran, Kaushik Lakshminarayanan, and Scott Shenker. *Resolving Inter-Domain Policy Disputes*. ACM SIGCOMM, August 2007.
- Hakim Weatherspoon, Patrick Eaton, **Byung-Gon Chun**, and John Kubiawicz. *Antiquity: Exploiting a Secure Log for Wide-Area Storage*. Second European Conference on Computer Systems (EuroSys '07), March 2007.
- **Byung-Gon Chun**, Frank Dabek, Andreas Haeberlen, Emil Sit, Hakim Weatherspoon, M. Frans Kaashoek, John Kubiawicz, and Robert Morris. *Efficient Replica Maintenance for Distributed Storage Systems*. Third USENIX/ACM Symposium on Networked Systems Design and Implementation (NSDI '06), May 2006.
- **Byung-Gon Chun**, Peter Wu, Hakim Weatherspoon, and John Kubiawicz. *ChunkCast: An Anycast Service for Large Content Distribution*. Fifth International Workshop on Peer-to-Peer Systems (IPTPS '06), February 2006.
- Emil Sit, Andreas Haeberlen, Frank Dabek, **Byung-Gon Chun**, Hakim Weatherspoon, Robert Morris, M. Frans Kaashoek, and John Kubiawicz. *Proactive Replication for Data Durability*. Fifth International Workshop on Peer-to-Peer Systems (IPTPS '06), February 2006.
- Sean Rhea, **Byung-Gon Chun**, John Kubiawicz, and Scott Shenker. *Fixing the Embarrassing Slowness of OpenDHT on PlanetLab*. Second Workshop on Real, Large Distributed Systems (WORLDS '05), December 2005. *Awarded Best Paper*.
- **Byung-Gon Chun**, Ben Y. Zhao, and John Kubiawicz. *Impact of Neighbor Selection on Performance and Resilience of Structured P2P Networks*. Fourth International Workshop on Peer-to-Peer Systems (IPTPS '05), February 2005.
- **Byung-Gon Chun**, Kamalika Chaudhuri, Hoeteck Wee, Marco Barreno, Christos H. Papadimitriou, and John Kubiawicz. *Selfish Caching in Distributed Systems: A Game-Theoretic Analysis*. Twenty-Third ACM Symposium on Principles of Distributed Computing (PODC '04), July 2004.
- **Byung-Gon Chun**, Rodrigo Fonseca, Ion Stoica, and John Kubiawicz. *Characterizing Selfishly Constructed Overlay Routing Networks*. Twenty-Third IEEE Conference on Computer Communications (INFOCOM '04), March 2004.
- **Byung-Gon Chun** and Mary Baker. *Evaluation of Packet Scheduling Algorithms in Mobile Ad Hoc Networks*. ACM Mobile Computing and Communications Review, Vol. 6, Issue 3, July 2002.

TECHNICAL REPORTS

- **Byung-Gon Chun**, Kuang Chen, Gunho Lee, Randy Katz, and Scott Shenker. *D3: Declarative Distributed Debugging*. U.C. Berkeley Technical Report No. EECS-2008-27, March 2008.
- Cheng Tien Ee, Vijay Ramachandran, **Byung-Gon Chun**, and Scott Shenker. *Resolving BGP Disputes*. U.C. Berkeley Technical Report No. EECS-2006-39, April 2006.
- Sriram Sankararaman, **Byung-Gon Chun**, Chawathe Yatin, and Scott Shenker. *Key Consistency in DHTs*. U.C. Berkeley Technical Report UCB/EECS-2005-21, November 2005.

- Hakim Weatherspoon, **Byung-Gon Chun**, Chiu Wah So, and John Kubiatoicz. *Long-term Data Maintenance in Wide-area Storage Systems: A Quantitative Approach*. U.C. Berkeley Technical Report UCB/EECS-05-1404, July 2005.

REFERENCES

Prof. Scott Shenker
University of California, Berkeley
Computer Science Division
RAD Lab – MS 1776, 465 Soda Hall
Berkeley, CA 94720
Phone: (510) 643-3403
E-mail: shenker@icsi.berkeley.edu

Prof. John Kubiatoicz
University of California, Berkeley
Computer Science Division
673 Soda Hall
Berkeley, CA 94720
Phone: (510) 643-6817
E-mail: kubitron@cs.berkeley.edu

Dr. Petros Maniatis
Intel Research Berkeley
2150 Shattuck Avenue, Suite 1300
Berkeley, CA 94704
Phone: (510) 495-3087
E-mail: petros.maniatis@intel.com

Dr. Sylvia Ratnasamy
Intel Research Berkeley
2150 Shattuck Avenue, Suite 1300
Berkeley, CA 94704
Phone: (510) 495-3043
E-mail: sylvia.p.ratnasamy@intel.com