

Haifeng Yu

Intel Research Pittsburgh / Carnegie Mellon University
Pittsburgh, PA 15213, U.S.A.

Phone: (412)-427-6877 **Email:** yhf@cs.cmu.edu **URL:** <http://www.cs.cmu.edu/~yhf>

Research Interests

Distributed Systems, Fault-tolerance, Peer-to-Peer Computing, Database Systems, and Operating Systems.

Education

- Ph.D. in Computer Science, Duke University, September 2002. Thesis Advisor: Amin Vahdat
Thesis: "Wide-Area Replication Using Continuous Consistency: Theory and Practice"
Nominated for ACM Doctoral Dissertation Award
- M.S. in Computer Science, Duke University, December 1999. Thesis Advisor: Gershon Kedem
Thesis: "DRAM-Page Based Prediction and Prefetching"
- B.E. in Computer Science, Shanghai Jiao Tong University, P.R. China, June 1997.
Graduated with Honors

Experience

- Adjunct Assistant Professor (January 2004 to present), Computer Science Department, Carnegie Mellon University.
- Researcher (November 2003 to present), Intel Research Pittsburgh.
- Post-doctoral Researcher (September 2002 to November 2003), Computer Science Department, Duke University.

Refereed Publications

Available online at <http://www.cs.cmu.edu/~yhf>

Following conferences in the publication list are often considered first-tier in their respective areas: *SOSP*, *OSDI*, *NSDI* (Systems); *VLDB* (Database); *PODC* (Distributed computing); *SIGCOMM* (Networking).

1. **Haifeng Yu**, Michael Kaminsky, Phillip B. Gibbons, and Abraham Flaxman, "Defending against Sybil Attacks via Social Networks", *Proceedings of ACM SIGCOMM Conference*, September 2006.
2. **Haifeng Yu**, Phillip B. Gibbons, and Suman Nath, "Availability of Multi-Object Operations", *Proceedings of the Symposium on Networked Systems Design and Implementation (NSDI)*, May 2006. **Awarded Best Paper**. (Acceptance rate: 28 out of 110)
3. Suman Nath, **Haifeng Yu**, Phillip B. Gibbons, and Srinivasan Seshan, "Subtleties in Tolerating Correlated Failures", *Proceedings of the Symposium on Networked Systems Design and Implementation (NSDI)*, May 2006. (Acceptance rate: 28 out of 110)
4. Michael Freedman, Scott Garriss, Michael Kaminsky, Brad Karp, David Mazieres, and **Haifeng Yu**, "RE: Reliable Email", *Proceedings of the Symposium on Networked Systems Design and Implementation (NSDI)*, May 2006. (Acceptance rate: 28 out of 110)
5. **Haifeng Yu**, "Signed Quorum Systems", *Distributed Computing*, Volume 18, Number 4 (*Special Journal Issue for PODC-04*), March 2006.

6. **Haifeng Yu** and Amin Vahdat, "The Costs and Limits of Availability for Replicated Services", *ACM Transactions on Computer Systems (TOCS)*, Volume 24, Issue 1, February 2006.
7. **Haifeng Yu** and Amin Vahdat, "Consistent and Automatic Replica Regeneration", *ACM Transactions on Storage (TOS)*, Volume 1, Issue 1, February 2005.
8. Praveen Yalagandula, Suman Nath, **Haifeng Yu**, Phillip Gibbons, and Srinivasan Seshan, "Beyond Availability: Towards a Deeper Understanding of Machine Failure Characteristics in Large Distributed Systems", *Proceedings of the Workshop on Real, Large Distributed Systems (WORLDS)*, December 2004.
9. **Haifeng Yu**, "Signed Quorum Systems", *Proceedings of the 23rd ACM Symposium on Principles of Distributed Computing (PODC)*, July 2004. (Acceptance rate: 39 out of 224)
10. **Haifeng Yu** and Amin Vahdat, "Consistent and Automatic Replica Regeneration", *Proceedings of the Symposium on Networked Systems Design and Implementation (NSDI)*, March 2004. (Acceptance rate: 27 out of 118)
11. **Haifeng Yu**, "Overcoming the Majority Barrier in Large-Scale Systems", *Proceedings of the 17th International Symposium on Distributed Computing (DISC)*, October 2003. (Acceptance rate: 25 out of 90)
12. Roger Barga, David Lomet, Stelios Paparizos, **Haifeng Yu**, and Sirish Chandrasekaran, "Persistent Applications via Automatic Recovery", *Proceedings of the 17th International Database Engineering and Applications Symposium (IDEAS)*, July 2003.
13. **Haifeng Yu** and Amin Vahdat, "Design and Evaluation of a Conit-based Continuous Consistency Model for Replicated Services", *ACM Transactions on Computer Systems (TOCS)*, Volume 20, Issue 3, August 2002.
14. **Haifeng Yu** and Amin Vahdat, "Minimal Replication Cost for Availability", *Proceedings of the 21st ACM Symposium on Principles of Distributed Computing (PODC)*, July 2002.
15. **Haifeng Yu** and Amin Vahdat, "The Costs and Limits of Availability for Replicated Services", *Proceedings of the 18th Symposium on Operating Systems Principles (SOSP)*, October 2001. (Acceptance rate: 17 out of 85)
16. **Haifeng Yu** and Amin Vahdat, "Combining Generality and Practicality in a Conit-Based Continuous Consistency Model for Wide-Area Replication", *Proceedings of the 21st International Conference on Distributed Computing Systems (ICDCS)*, April 2001. (Acceptance rate: 72 out of 212)
17. **Haifeng Yu** and Amin Vahdat, "Design and Evaluation of a Continuous Consistency Model for Replicated Services", *Proceedings of the Fourth Symposium on Operating Systems Design and Implementation (OSDI)*, October 2000. (Acceptance rate: 24 out of 111)
18. **Haifeng Yu** and Amin Vahdat, "Efficient Numerical Error Bounding for Replicated Network Services", *Proceedings of the 26th International Conference on Very Large Databases (VLDB)*, September 2000. (Acceptance rate: 53 out of 351)

19. **Haifeng Yu** and Gershon Kedem, "DRAM-Page Based Prediction and Prefetching", *Proceedings of the International Conference on Computer Design (ICCD)*, September 2000.
20. **Haifeng Yu** and Amin Vahdat, "Building Replicated Internet Services Using TACT: A Toolkit for Tunable Availability and Consistency Tradeoffs", *Proceedings of the Second International Workshop on Advanced Issues of E-Commerce and Web-based Information Systems (WECWIS)*, June 2000.

Reference / Citation Count

- Over 200 citations: See <http://scholar.google.com>
- Papers have been used as course material in many universities, including Cornell University, University of Rochester, Georgia Institute of Technology, University of Maryland (College Park), Rutgers University, University of Massachusetts (Amherst), Northwestern University, University of Colorado (Boulder), and University of Texas (Austin).

Conference Program Committee Member of

- The 1st ACM Workshop on Scalable Trusted Computing (STC), 2006.
- The 25th ACM Symposium on Principles of Distributed Computing (PODC), 2006.
- The International Conference on Dependable Systems and Networks (DSN), 2006.
- The 24th ACM Symposium on Principles of Distributed Computing (PODC), 2005.
- External reviewer for NSDI, ICDCS, Usenix Technical Conference, HotOS, ASPLOS, Usenix Symposium on Internet Technologies, FAST, and others.

Journal Referee for

- ACM Transactions on Storage (ToS)
- IEEE/ACM Transactions on Networking (ToN)
- IEEE Transactions on Dependable and Secure Computing (TDSC)
- IEEE Transactions on Mobile Computing (TMC)
- IEEE Transactions on Parallel and Distributed Systems (TPDS)
- IEEE Transactions on Systems, Man, and Cybernetics, Part A (TSMC)
- Journal of Parallel and Distributed Computing (JPDC)
- World Wide Web Journal (WWWJ)

Invited Talks / Conference Talks

1. "Robust Distributed Systems despite Failure Correlation", *invited talk at UC San Diego*, November 2005.
2. "Availability of Multi-Object Operations", *invited talk at HP Labs*, December 2004.
3. "Availability of Multi-Object Operations", *invited talk at Intel Research Berkeley*, December 2004.
4. "Consistent and Automatic Replica Regeneration", *Symposium on Networked Systems Design and Implementation (NSDI)*, March 2004.
5. "Consistent and Automatic Replica Regeneration", *invited talk at UT Austin*, March 2004.
6. "Consistent and Automatic Replica Regeneration", *invited talk at UC San Diego*, March 2004.

7. "Minimal Replication Cost for Availability", *21st Symposium on Principles of Distributed Computing (PODC)*, July 2002.
8. "Wide-Area Replication Using Continuous Consistency: Theory and Practice", *invited talk at MIT*, June 2002.
9. "TACT: Continuous Consistency for Wide-Area Replication", *invited talk at University of Texas, Austin*, February 2002.
10. "The Costs and Limits of Availability for Replicated Services", *18th Symposium on Operating Systems Principles (SOSP)*, October 2001.
11. "TACT: Continuous Consistency for Wide-Area Replication", *invited talk at Microsoft Research and University of Washington, Seattle*, August 2001.
12. "Combining Generality and Practicality in a Conit-Based Continuous Consistency Model for Wide-Area Replication", *21st International Conference on Distributed Computing Systems*, April 2001.
13. "Building Replicated Internet Services Using TACT: A Toolkit for Tunable Availability and Consistency Tradeoffs", *2nd International Workshop on Advanced Issues of E-Commerce and Web-based Information Systems*, June 2000.