

Haifeng YU

Department of Computer Science, National University of Singapore
COM2-04-25, 15 Computing Drive, Republic of Singapore 117418

haifeng@comp.nus.edu.sg
<http://www.comp.nus.edu.sg/~yuhf>

July 2011

Research Interests

Distributed Computing, Distributed Algorithms, Distributed Systems Security, Distributed Systems Availability, and Applied Algorithms in Networking.

Professional Experience

- **July 2010 — present:** Associate Professor (with Tenure), School of Computing, National University of Singapore, Republic of Singapore.
- **Oct 2006 — June 2010:** Assistant Professor, School of Computing, National University of Singapore, Republic of Singapore.
- **Nov 2003 — Sept 2006:** Research Scientist, Intel Research, Pittsburgh, USA.
- **Nov 2003 — Sept 2006:** Adjunct Assistant Professor, Computer Science Department, Carnegie Mellon University, Pittsburgh, USA.
- **Sept 2002 — Nov 2003:** Post-doctoral Researcher, Computer Science Department, Duke University, USA.

Education

- **Aug 1997 — Sept 2002:** Ph.D. in Computer Science, Duke University, USA.
Thesis Advisor: Amin Vahdat. Thesis: "*Wide-Area Replication Using Continuous Consistency: Theory and Practice*". Nominated for ACM Doctoral Dissertation Award.
- **Aug 1997 — Dec 1999:** M.S. in Computer Science, Duke University, USA.
Thesis Advisor: Gershon Kedem. Thesis: "*DRAM-Page Based Prediction and Prefetching*"
- **May 1993 — June 1997:** B.E. in Computer Science, Shanghai Jiao Tong University, P.R. China.
Thesis: "*Thread-based Fault-tolerant Distributed Shared Memory*"

Research Awards

- **Best Paper Award** in ACM SIGCOMM Conference (*SIGCOMM'10*), 2010.
- **Best Paper Award** in ACM/IEEE International Conference on Information Processing in Sensor Networks (*IPSN'09*), 2009.
- **Best Paper Award** in Symposium on Networked Systems Design & Implementation (*NSDI'06*), 2006.
- **Young Scientist Award 2009 from Singapore National Academy of Science:** This is a national award given out by Singapore National Academy of Science, "to young researchers, aged 35 years and below, who are actively engaged in R&D in Singapore, and who have shown great potential to be world-class researchers in their fields of expertise." Since the awards launch in 1997 and until 2009, the award was given each year to **no more than 3 winners, across all disciplines of natural sciences**. The award came with a cash prize (not research grant) of S\$10,000 (about USD\$7,000). There was

also media coverage in Singapore's main newspaper. I won this award for my "work on defending against sybil attacks and on distributed systems security".

- **Young Researcher Award 2011 from National University of Singapore:** This award was "conferred to researchers below 40 years of age based on their impact and promise in research." In 2011, the award was given to 4 winners, across all departments in NUS. The award came with a cash prize (not research grant) of S\$2,000 (about USD\$1,500).
- **Young Investigator Award 2007 from National University of Singapore:** "The NUS Young Investigator Award is an initiative to offer research-funding to highly-promising young faculty members." In 2007, the award was given to 5 winners, across all departments in NUS. My research proposal "Towards Secure and Highly-Available Aggregation Queries in Large-Scale Sensor Networks" won this award, with a total project value of S\$426,602 (about USD\$300,000).

Some Potential Research Impact Indicators

- **Citation count from Google Scholar:** Over 1,500 as of July 2011.
- **Used as course materials:** My publications have been used as course materials in graduate-level courses over 50 times in major universities worldwide, such as UC Berkeley, Columbia University, Cornell University, EPFL, etc. Some of them have also been included as an entire section in a popular graduate-level textbook on distributed systems. One of my publications has been included in the "Basic Reading List" of the Ph.D. Qualifying Exam in UIUC. (As of June 2010.)
- **Follow-up research:** *Follow-up research* here is narrowly defined as research work done by other people that adopts/adapts the novel observations, insights, or techniques developed in my work to solve related problems. Research work that simply uses (as a black box) the end results or conclusions of my work is not considered as follow-up research. My research has been followed-up by people from high-profile places such as MIT, New York University, Microsoft Research, Rice University, University of Texas at Austin, Brown University, University of Utah, University of Maryland, and etc., with follow-up publications in high-quality publication venues such as SIGCOMM, NDSS, and NSDI.

Editorial Board Membership

- International Journal of Security and Networks (IJSN)

Program Committee Member

- PC co-Chair for Distributed Computing Track, ICDCN'11 (International Conference on Distributed Computing and Networking)
- SIGCOMM'12 (ACM SIGCOMM Conference)
- APSys'11 (ACM SIGOPS Asia-Pacific Workshop on Systems)
- CCS'11 (ACM Conference on Computer and Communications Security)
- SIGCOMM'11 (ACM SIGCOMM Conference)
- SIGMETRICS'11 (International Conference on Measurement and Modeling of Computer Systems)
- CCS'10 (ACM Conference on Computer and Communications Security)
- Oakland'10 (IEEE Symposium on Security and Privacy)
- IPTPS'10 (International Workshop on Peer-to-Peer Systems)
- IPDPS'10 (IEEE International Parallel and Distributed Processing Symposium)

- ICDCN'10 (International Conference on Distributed Computing and Networking)
- NetEcon'10 (Workshop on the Economics of Networks, Systems, and Computation)
- DCOSS'10 (IEEE International Conference on Distributed Computing in Sensor Systems)
- PODC'09 (ACM Symposium on Principles of Distributed Computing)
- ICDCS'09 (International Conference on Distributed Computing Systems)
- ICDE'09 (International Conference on Data Engineering)
- DISC'08 (International Symposium on Distributed Computing)
- SIGMETRICS'08 (International Conference on Measurement and Modeling of Computer Systems)
- DSN'08 (International Conference on Dependable Systems and Networks)
- STC'07 (ACM Workshop on Scalable Trusted Computing)
- IPTPS'07 (International Workshop on Peer-to-Peer Systems)
- PODC'06 (ACM Symposium on Principles of Distributed Computing)
- DSN'06 (International Conference on Dependable Systems and Networks)
- STC'06 (ACM Workshop on Scalable Trusted Computing)
- PODC'05 (ACM Symposium on Principles of Distributed Computing)

Invited Papers

- [I1] **Haifeng Yu.**
 Sybil Defenses via Social Networks: A Tutorial and Survey.
 ACM SIGACT News Distributed Computing Column, to appear.

Journal Publications

- [J1] Binbin Chen, Ziling Zhou, Yuda Zhao, and **Haifeng Yu.**
 Efficient Error Estimating Coding: Feasibility and Applications.
 IEEE/ACM Transactions on Networking (ToN), to appear.
- [J2] **Haifeng Yu.**
 Secure and Highly-Available Aggregation Queries in Large-Scale Sensor Networks via Set Sampling.
 Distributed Computing, Volume 23, Issue 5, April 2011.
- [J3] **Haifeng Yu**, Phillip B. Gibbons, Michael Kaminsky, and Feng Xiao.
 SybilLimit: A Near-Optimal Social Network Defense against Sybil Attacks.
 IEEE/ACM Transactions on Networking(ToN), Volume 18, Issue 3, June 2010.
- [J4] **Haifeng Yu** and Phillip B. Gibbons.
 Optimal Inter-Object Correlation When Replicating for Availability.
 Distributed Computing, Volume 21, Number 5, February 2009. **Special journal issue for PODC'07.**
- [J5] **Haifeng Yu**, Michael Kaminsky, Phillip B. Gibbons, and Abraham Flaxman.
 SybilGuard: Defending Against Sybil Attacks via Social Networks.
 IEEE/ACM Transactions on Networking (ToN), Volume 16, Issue 3, June 2008.
- [J6] **Haifeng Yu.**
 Signed Quorum Systems.
 Distributed Computing, Volume 18, Number 4, March 2006. **Special journal issue for PODC'04.**
- [J7] **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.

- [J8] **Haifeng Yu** and Amin Vahdat.
Consistent and Automatic Replica Regeneration.
ACM Transactions on Storage (TOS), Volume 1, Number 1, February 2005.
- [J9] **Haifeng Yu** and Amin Vahdat.
Design and Evaluation of a Conit-based Continuous Consistency Model for Replicated Services.
ACM Transactions on Computer Systems (TOCS), August 2002.

Major Conference Publications

- [C1] Binbin Chen, Ziling Zhou, Yuda Zhao, and **Haifeng Yu**.
Efficient Error Estimating Coding: Feasibility and Applications.
In Proceedings of ACM SIGCOMM Conference (*SIGCOMM'10*), August 2010. **Awarded Best Paper.**
- [C2] Suman Nath, **Haifeng Yu**, and Haowen Chan.
Secure Outsourced Aggregation via One-way Chains.
In Proceedings of the ACM SIGMOD Conference (*SIGMOD'09*), June 2009.
- [C3] **Haifeng Yu**, Chenwei Shi, Michael Kaminsky, Phillip B. Gibbons, and Feng Xiao.
DSybil: Optimal Sybil-Resistance for Recommendation Systems.
In Proceedings of the IEEE Symposium on Security and Privacy (*Oakland'09*), May 2009.
- [C4] **Haifeng Yu**.
Secure and Highly-Available Aggregation Queries in Large-Scale Sensor Networks via Set Sampling.
In Proceedings of the ACM/IEEE International Conference on Information Processing in Sensor Networks (*IPSN'09*), April 2009. **Awarded Best Paper.**
- [C5] **Haifeng Yu**, Phillip B. Gibbons, Michael Kaminsky, and Feng Xiao.
SybilLimit: A Near-Optimal Social Network Defense against Sybil Attacks.
In Proceedings of the IEEE Symposium on Security and Privacy (*Oakland'08*), May 2008.
- [C6] **Haifeng Yu** and Phillip B. Gibbons.
Optimal Inter-Object Correlation When Replicating for Availability.
In Proceedings of the ACM Symposium on Principles of Distributed Computing (*PODC'07*), August 2007. **Selected to be fast-tracked to *Distributed Computing*.**
- [C7] **Haifeng Yu**, Michael Kaminsky, Phillip B. Gibbons, and Abraham Flaxman.
SybilGuard: Defending Against Sybil Attacks via Social Networks.
In Proceedings of ACM SIGCOMM Conference (*SIGCOMM'06*), September 2006. **Selected to be fast-tracked to *Transactions on Networking*.**
- [C8] **Haifeng Yu**, Phillip B. Gibbons, and Suman Nath.
Availability of Multi-Object Operations.
In Proceedings of the Symposium on Networked Systems Design and Implementation (*NSDI'06*), May 2006. **Awarded Best Paper.**
- [C9] Suman Nath, **Haifeng Yu**, Phillip B. Gibbons, and Srinivasan Seshan.
Subtleties in Tolerating Correlated Failures in Wide-area Storage Systems.
In Proceedings of the Symposium on Networked Systems Design and Implementation (*NSDI'06*), May 2006.
- [C10] **Haifeng Yu**.
Signed Quorum Systems.
In Proceedings of the ACM Symposium on Principles of Distributed Computing (*PODC'04*), July 2004. **Selected to be fast-tracked to *Distributed Computing*.**
- [C11] **Haifeng Yu** and Amin Vahdat.
Consistent and Automatic Replica Regeneration.
In Proceedings of the Symposium on Networked Systems Design and Implementation (*NSDI'04*), March 2004.

- [C12] **Haifeng Yu.**
Overcoming the Majority Barrier in Large-Scale Systems.
In Proceedings of the International Symposium on Distributed Computing (*DISC'03*), October 2003.
- [C13] **Haifeng Yu** and Amin Vahdat.
Minimal Replication Cost for Availability.
In Proceedings of the ACM Symposium on Principles of Distributed Computing (*PODC'02*), July 2002.
- [C14] **Haifeng Yu** and Amin Vahdat.
The Costs and Limits of Availability for Replicated Services.
In Proceedings of the Symposium on Operating Systems Principles (*SOSP'01*), October 2001.
- [C15] **Haifeng Yu** and Amin Vahdat.
Design and Evaluation of a Continuous Consistency Model for Replicated Services.
In Proceedings of the Symposium on Operating Systems Design and Implementation (*OSDI'00*), October 2000.
- [C16] **Haifeng Yu** and Amin Vahdat.
Efficient Numerical Error Bounding for Replicated Network Services.
In Proceedings of the International Conference on Very Large Databases (*VLDB'00*), September 2000.

Other Publications

- [O1] Binbin Chen and **Haifeng Yu.**
Secure Aggregation with Malicious Node Revocation in Sensor Networks.
In Proceedings of International Conference on Distributed Computing Systems (*ICDCS'11*), June 2011.
- [O2] **Haifeng Yu**, Phillip B. Gibbons, and Chenwei Shi.
Brief Announcement: Sustaining Collaboration in Multicast despite Rational Collusion.
In Proceedings of the ACM Symposium on Principles of Distributed Computing (*PODC'11*), June 2011.
- [O3] **Haifeng Yu**, Phillip B. Gibbons, and Michael Kaminsky.
Brief Announcement: Toward an Optimal Social Network Defense against Sybil Attacks.
In Proceedings of the ACM Symposium on Principles of Distributed Computing (*PODC'07*), August 2007.
- [O4] **Haifeng Yu.**
Brief Announcement: DoS-Resilient Secure Aggregation Queries in Sensor Networks.
In Proceedings of the ACM Symposium on Principles of Distributed Computing (*PODC'07*), August 2007.
- [O5] Jeffrey Pang, Phillip B. Gibbons, Michael Kaminsky, Srinivasan Seshan, and **Haifeng Yu.**
Defragmenting DHT-based Distributed File Systems.
Proceedings of International Conference on Distributed Computing Systems (*ICDCS'07*), June 2007.
- [O6] Amit Manjhi, Phillip B. Gibbons, Anastassia Ailamaki, Charles Garrod, Bruce M. Maggs, Todd C. Mowry, Christopher Olston, Anthony Tomasic, and **Haifeng Yu.**
Invalidation Clues for Database Scalability Services.
In Proceedings of International Conference on Data Engineering (*ICDE'07*), April 2007.
- [O7] Scott Garriss, Michael Kaminsky, Michael Freedman, Brad Karp, David Mazieres, and **Haifeng Yu.**
RE: Reliable Email.
In Proceedings of the Symposium on Networked Systems Design and Implementation (*NSDI'06*), May 2006.
- [O8] Praveen Yalagandula, Suman Nath, **Haifeng Yu**, Phillip B. Gibbons, and Srinivasan Seshan.
Beyond Availability: Towards a Deeper Understanding of Machine Failure Characteristics in Large Distributed Systems.
In Proceedings of the Workshop on Real, Large Distributed Systems (*WORLDS'04*), December 2004.

- [O9] Roger Barga, David Lomet, Stelios Paparizos, **Haifeng Yu**, and Sirish Chandrasekaran.
Persistent Applications via Automatic Recovery.
In Proceedings of the International Database Engineering and Applications Symposium (*IDEAS'03*), July 2003.
- [O10] **Haifeng Yu** and Amin Vahdat.
Combining Generality and Practicality in a Conit-Based Continuous Consistency Model for Wide-Area Replication.
In Proceedings of the International Conference on Distributed Computing Systems (*ICDCS'01*), April 2001.
- [O11] **Haifeng Yu** and Gershon Kedem.
DRAM-Page Based Prediction and Prefetching.
In Proceedings of International Conference on Computer Design (*ICCD'00*), September 2000.
- [O12] **Haifeng Yu** and Amin Vahdat.
Building Replicated Internet Services Using TACT: A Toolkit for Tunable Availability and Consistency Tradeoffs.
In Proceedings of the International Workshop on Advanced Issues of E-Commerce and Web-based Information Systems (*WECWIS'00*), June 2000.

Teaching Experience in National University of Singapore

- AY 2011/2012 Semester 2: CS3230 Design and Analysis of Algorithms
- AY 2011/2012 Semester 1: CS5223 Distributed Systems
- AY 2010/2011 Semester 2: CS4231 Parallel and Distributed Algorithms
- AY 2010/2011 Semester 1: CS5223 Distributed Systems
- AY 2009/2010 Semester 2: CS3230 Design and Analysis of Algorithms
- AY 2009/2010 Semester 1: CS5223 Distributed Systems
- AY 2008/2009 Semester 2: CS4231 Parallel and Distributed Algorithms
- AY 2008/2009 Semester 1: CS5223 Distributed Systems
- AY 2007/2008 Semester 2: CS5321 Network Security and Management
- AY 2007/2008 Semester 1: CS5223 Distributed Systems
- AY 2006/2007 Semester 2: CS4231 Parallel and Distributed Algorithms

Research Grants in National University of Singapore

Note: Currently my university provides centrally allocated scholarships for Ph.D. students, so there is usually no need to use grant money to pay for Ph.D. students working with me. Also given the algorithmic flavor of my research, my research work does not usually entail large equipment budgets.

I am the sole PI on all the following grants. None of the grants has a co-PI.

- **Start-up grant:** “Meeting New Challenges in Highly Available Decentralized Distributed Systems”, S\$173,716 (about USD\$124,000), Feb 2007 to Jul 2009.
- **NUS Young Investigator Award:** “Towards Secure and Highly-Available Aggregation Queries in Large-Scale Sensor Networks”, S\$426,602 (about USD\$300,000), Jan 2008 to July 2011.
- **NUS FRC grant:** “Combating Selfish and Malicious Behaviour in Large-Scale Decentralized Distributed Systems”, S\$69,460 (about USD\$50,000), Oct 2009 to Sep 2012.

People (Being) Supervised in National University of Singapore

- Binbin CHEN (Post-doctoral researcher, now at ADSC)
- Yuda ZHAO, Liu Xiao (Ph.D. students)
- Chenwei SHI, Feng XIAO (Master student)
- Wensi WANG, Haojun ZHANG, Yong WANG, Xiaojiang XU (Undergraduate students)