

Fan Ye

Contact Information

1870 Baldwin Road, Apt 57
Yorktown Heights, NY 10598, USA
Email: yefann@gmail.com
Homepage: <http://www.research.ibm.com/people/y/yefan/>

Phone: (917) 340 6172 (C)
(914) 784 7299 (O)

Education

Ph.D. (Networking), Computer Science, UCLA 2004.6, GPA 4.0/4.0
Thesis: “Building Resilient Sensor Networks by Exploiting Scale and Redundancy”
Advisor: Prof. Lixia Zhang
M.S., Computer Science, Tsinghua University, China 1999.6, GPA 3.7/4.0
Thesis: “Design and Implementation of a Secure Web Server”
Advisor: Prof. Daoyuan Hu
B.E., Automation, Tsinghua University, China 1996.6, GPA 3.8/4.0

Research Experience

IBM T.J. Watson Research Center Research Staff Member, full time, 40 hours/week
Next Generation Distributed Systems Dept. August 2004 – present
I have worked on several areas in large scale streaming processing, including federated resource discovery, data source management and cooperative stream processing. Scalable stream processing has wide applications in decision support for business, government and science. The next generation stream processing system that I work on can scale magnitudes higher than existing systems. It allows many users to pose semi-natural language questions and analyzes huge amount of streaming data to give answers. I have also worked on secure traceback of colluding insider attacks in sensor networks.

- **Federated Resource Discovery** allows stream processing applications to locate and utilize resources that are distributed over wide area networks. It supports a unified search interface that lets applications find resources in many different autonomous systems, thus utilizing resources from all of them.
- **Data Source Management** is the focal point in a stream processing system to discover and allocate data sources for applications. It maintains the characteristic descriptions of potentially millions of data sources, and select the most suitable ones given applications’ requirements. It also monitors the dynamic changes of both external and internal data sources.
- **Cooperative Stream Processing** enables multiple autonomous stream processing systems to collaborate and further scale up the processing capabilities. It allows analysis and processing that are difficult or impossible in any single system. An architecture that supports flexible collaboration models is the key for systems belonging to different organizations to work together.
- **Traceback in Sensor Networks** allows users to identify the exact origin of attacking packets, such that compromised insider nodes that inject bogus data can be isolated from the network. This is the first work that addresses colluding attacks where multiple compromised nodes work together to cover their traces.

Internet Research Laboratory

Research Assistant, part-time, 20 hours/week

Computer Science Dept., UCLA

September 1999 – present

Sponsored by the DARPA SensIT program, I have worked on *data forwarding, energy efficiency and security* in sensor networks. I designed, analyzed and evaluated protocols that build resilient, large-scale sensor networks operating robustly under severe conditions. My papers have been used as course materials in schools including Brown, UCLA, Boston University, University of Massachusetts, etc.

Data Forwarding

- **GRADient Broadcast (GRAB)**: GRAB utilizes a novel concept of cost-field to enable *robust* and *scalable* data forwarding under extremely severe conditions, which is critical for sensor networks to work in adverse or hostile environments. Instead of maintaining state about explicit data paths, nodes make forwarding decisions based on state carried in packets. Thus data delivery is immune to significant node failures and channel errors. Publications from this project have been **cited 278 times**¹
- **Two-tier Data Dissemination (TTDD)**: TTDD is the first protocol that delivers data efficiently to multiple users that are in *constant motion*, which is a typical scenario for many sensor network applications. Different from “flat” routing where every node plays the same role (as in most sensor network forwarding protocols), it builds a virtual hierarchy to stabilize data forwarding routes. A user’s mobility affects data forwarding in the local area but becomes transparent to the global route, which remains relatively stable. Publications from this project have been **cited 335 times**.

Energy Efficiency

- **Probing Environment and Adaptive Sleeping (PEAS)**: PEAS extends system lifetime in proportion to nodes’ deployment density. It controls the density of working nodes by keeping a certain distance between working neighbors, like trees in a forest. This leads to a simple, yet robust and scalable design; the only information nodes need is whether any working neighbor exists within a certain range. Publications in this projects have been **cited 366 times**.

Security

- **Statistical En-route Filtering (SEF)**: SEF is the first proposal that detects bogus data injected by *compromised insider* nodes. It enables the network to filter such data *en-route*, thus avoiding grave consequences in mission-critical applications, such as wrong decisions be made in battlefields, energy exhausted in networks. It assigns secret keys to nodes in a carefully designed way, thus each forwarding node can verify the truthfulness of data packets probabilistically. By accumulating the detecting power of nodes, most false data can be dropped before reaching the destination. Publications from this project have been **cited 127 times**.
- **Location-Dependent Keys (LDK)**: LDK provides secure report generation that rejects bogus data and ensures the reporting of legitimate data against *large numbers* of compromised nodes. In LDK, keys a node possesses are not decided by its identity (as in other existing work), but by its location. This binding between keys and locations enables the user to verify if a report does originate from the claimed location. Publications from this work have been **cited 38 times**.

I also participated in a project for recursive localization of sensor networks. More details of my research are available at <http://www.research.ibm.com/people/y/yefan/>.

Network Research Laboratory

Computer Science Dept, Tsinghua Univ. China

Research Assistant

September 1996 – June 1999

I designed and implemented a KDC-based secure web server and its client-side proxy. They work seamlessly with legacy web browser and server software to provide user authentication, fine-grain access control and communication privacy.

Teaching Experience

UCLA Computer Science Dept.

Supervised 3 undergraduate and junior graduate students on research projects (GRAB, PEAS, TTDD).

Guest lecture to a graduate course (CS 211).

ExecuTrain

Beijing, China

Certified Computer Training Instructor

November 1996 – September 1998

I was one of the first 6 ExecuTrain Certified Instructors in China and taught various technical courses. ExecuTrain is an Atlanta-based company leading in business technology training.

¹Citation numbers from <http://scholar.google.com/> as of 12/23/07

Awards

May 2007	IBM Research Invention Achievement Plateau Level 2
April 2007	IBM Research Invention Achievement Plateau Level 1
June 2006	IBM Research Invention Achievement Award
May 2003	ICDCS'03 travel award
April 2003	IPSN'03 travel award
July 1996	Prominent Graduate, Tsinghua University
June 1995	Motorola Prize, Tsinghua University
June 1994	Motorola Prize, Tsinghua University
June 1993	Excellent Student Second Prize, Tsinghua University
June 1992	Excellent Student Third Prize, Tsinghua University

Presentations

Conference and Invited Talk

1. "A Scalable Solution to Minimum Cost Forwarding in Large Sensor Networks", in the 10th International Conference on Computer Communications and Networks (IEEE ICCN), October 2001, Scottsdale, Arizona, USA
2. "PEAS: A Robust Energy Conserving Protocol for Long-lived Sensor Networks," in the 23rd International Conference on Distributed Computing Systems (IEEE ICDCS), May 2003, Providence, Rhode Island, USA
3. "Statistical En-route Detection and Filtering of Injected False Data in Sensor Networks," in the 23rd IEEE Conference on Computer Communications (IEEE INFOCOM), March 2004, Hong Kong, China
4. "A Shared Infrastructure for Cooperative Sensing Applications", in the Fourth IEEE International Conference on Cooperative Internet Computing, Hong Kong, China, October 2006
5. "Catching 'Moles' in Sensor Networks", invited talk, University of Connecticut, CSE department, September 2006
6. "Catching 'Moles' in Sensor Networks", IEEE International Conference on Distributed Computing Systems (ICDCS), Toronto, Canada, June 2007
7. "*CLASP: Collaborating, Autonomous Stream Processing Systems", ACM/USENIX International Middleware Conference, Newport Beach, California, November 2007

Poster Presentation

1. "Gradient Broadcast in Sensor Networks," in DARPA SensIT PI meeting, April 2001, St. Petersburg, Florida, USA.
2. "A Robust Data Delivery Protocol for Large Scale Sensor Networks," in the 2nd International Workshop on Information Processing in Sensor Networks (IPSN), April 2003, Palo Alto, California, USA.
3. "PEAS: A Robust Energy Conserving Protocol for Long-lived Sensor Networks," in the 10th IEEE International Conference on Network Protocols (IEEE ICNP), November 2002, Paris, France
4. "Statistical En-route Filtering in Large Scale Sensor Networks," the 1st International Conference on Embedded Networked Sensor Systems (extended abstract in ACM SenSys), November 2003, Los Angeles, California, USA.

Publications

² Refereed Journals, Book Chapters and Magazines

1. **Fan Ye**, Hao Yang, Zhen Liu, Starsky Wong, Songwu Lu, Lixia Zhang, Mani Srivastava, "Preserving Data Authenticity in Wireless Sensor Networks: Attacks and Countermeasures", invited book chapter in Wireless and Sensor Networks Security, 2007
2. **Fan Ye**, Honghai Zhang, Songwu Lu, Lixia Zhang, Jennifer Hou, "A Randomized Energy-Conservation Protocol for Resilient Sensor Networks", ACM Wireless Networks (WINET) Journal, Vol. 12, No. 5, pp 637-652, September 2006, **conference version cited by 301**

²* Authors ordered in alphabetic order for most IBM publications

3. Fred Douglis, Michael Branson, Kirsten Hildrum, Bin Rong, Fan Ye*, "Multi-site Cooperative Data Stream Analysis," ACM SIGOPS Operating Systems Review, Vol. 40, No. 3, pp 31-37, July 2006
4. **Fan Ye**, Haiyun Luo, Songwu Lu, Lixia Zhang, "Statistical En-route Filtering of Injected False Data in Sensor Networks," IEEE Journal on Selected Areas in Communications, Vol. 23, No. 4, pp 839-850, April, 2005, **cited by 127**
5. **Fan Ye**, Gary Zhong, Songwu Lu, Lixia Zhang, "GRAdient Broadcast: A Robust Data Delivery Protocol for Large Scale Sensor Networks," ACM Wireless Networks (WINET) Journal, pp 285-298, Vol. 11, No.3, May 2005, **cited by 127**
6. Haiyun Luo, **Fan Ye**, Jerry Cheng, Songwu Lu, Lixia Zhang, "TTDD: Two-tier Data Dissemination in Large-scale Sensor Networks", ACM Wireless Networks (WINET) Journal, Vol. 11, No. 1-2, January, 2005, **cited by 43**
7. **Fan Ye**, Haiyun Luo, Songwu Lu, Lixia Zhang, "Dissemination Protocols for Large Sensor Networks," book chapter in Wireless Sensor Networks, 2003 (invited submission).
8. Hao Yang, Haiyun Luo, Fan Ye, Jiejun Kong, Petros Zerfos, Songwu Lu, Lixia Zhang, Mario Gerla, "Security in Mobile and Ad Hoc Wireless Networks'," book chapter in Handbook of Mobile Computing, 2003 (invited submission).
9. Haiyun Luo, Hao Yang, Fan Ye, Songwu Lu, Lixia Zhang, "Security in Mobile Ad-Hoc Wireless Networks: Challenges and Solutions," in IEEE Wireless Communications Magazine, 2003 (invited submission), **cited by 109**
10. Ming Lu, **Fan Ye**, Daoyuan Hu, "A Research on WWW Security Mechanism," Electronic Science and Technology Review (in Chinese), January 1998.

Refereed Conferences and Workshops

1. Bo An, Fred Douglis, Fan Ye, "*Heuristics for Negotiation Schedules in Multi-plan Optimization," to appear in the Seventh International Joint Conference on Autonomous Agents and Multi-Agent Systems, Esotil, Portugal, May 2008 (acceptance ratio 22.1% 142/640)
2. Michael Branson, Fred Douglis, Brad Fawcett, Zhen Liu, Anton Riabov, Fan Ye, "*CLASP: Collaborating, Autonomous Stream Processing Systems", in ACM/USENIX International Middleware Conference, Newport Beach, California, November 2007 (acceptance ratio 20.6% 22/107)
3. Hao Yang, Fan Ye, Zhen Liu, "Resource Discovery in Federated Systems with Voluntary Sharing, in Work-in-Progress Track of ACM/USENIX International Middleware Conference, 2007
4. **Fan Ye**, Hao Yang, Zhen Liu, "Catching 'Moles' in Sensor Networks", IEEE International Conference on Distributed Computing Systems (ICDCS), Toronto, Canada, June 2007, (acceptance ratio 13.4% 71/528)
5. Michael Branson, Fred Douglis, Brad Fawcett, Zhen Liu, Anton Riabov, Fan Ye, "*Autonomic operations in cooperative stream processing systems", the IEEE/ACM Second Workshop on Hot Topics in Autonomic Computing (HotAC), Jacksonville, Florida, June, 2007
6. Eric Bouillet, Mark Feblowitz, Zhen Liu, Anand Ranganathan, Anton Riabov, Fan Ye, "*A Semantic Based Middleware for Utilizing Heterogeneous Sensor Networks", The 3rd IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS), Santa Fe, New Mexico, June 2007
7. Eric Bouillet, Mark Feblowitz, Zhen Liu, Anand Ranganathan, Anton Riabov, Schuman Shao, Don Schlosnagle, Fan Ye, "*Stream Processing Based Intelligent Transport Systems", The 7th International Conference on ITS Telecommunications, Sophia Antipolis, France, June 2007
8. Eric Bouillet, Mark Feblowitz, Zhen Liu, Anand Ranganathan, Anton Riabov, Schuman Shao, Don Schlosnagle, Fan Ye, "*Data Stream Processing Infrastructure for Intelligent Transport Systems", IEEE Vehicular Technology Conference, Fall, 2007

9. Zhen Liu, Dimitrios Pendarakis, Fan Ye, “*A Shared Infrastructure for Cooperative Sensing Applications”, the Fourth IEEE International Conference on Cooperative Internet Computing, Hong Kong, China, October 2006
10. Hao Yang, **Fan Ye**, Yuan Yuan, Songwu Lu, William Arbaugh, ”Toward Resilient Security in Wireless Sensor Networks”, ACM MobiHoc, pp 34-45, Urbana-Champaign, IL, May 2005, (acceptance ratio 14%, 40/281), **cited by 38**
11. **Fan Ye**, Haiyun Luo, Songwu Lu, Lixia Zhang, “Statistical En-route Detection and Filtering of Injected False Data in Sensor Networks,” in the 23rd IEEE Conference on Computer Communications (IEEE INFOCOM), March 2004, Hong Kong, China, (acceptance ratio 18%, 261/1420), **journal version cited by 127**
12. **Fan Ye**, Haiyun Luo, Jerry Cheng, Songwu Lu, Lixia Zhang, “ A Two-tier Data Dissemination Model for Large-scale Wireless Sensor Networks ”. in the 8th Annual International Conference on Mobile Computing and Networking (ACM MobiCom), September 2002, Atlanta, Georgia, USA. (acceptance ratio 7.1% 26/364), **cited by 292**
13. **Fan Ye**, Gary Zhong, Songwu Lu, Lixia Zhang, “PEAS: A Robust Energy Conserving Protocol for Long-lived Sensor Networks,” in the 23rd International Conference on Distributed Computing Systems (IEEE ICDCS), May 2003, Providence, Rhode Island, USA (acceptance ratio 17.7%, 72/406)
14. **Fan Ye**, Gary Zhong, Songwu Lu, Lixia Zhang, “A Robust Data Delivery Protocol for Large Scale Sensor Networks,” in the 2nd International Workshop on Information Processing in Sensor Networks (IPSN), April 2003, Palo Alto, California, USA.
15. **Fan Ye**, Haiyun Luo, Songwu Lu, Lixia Zhang, “Statistical En-route Filtering in Large Scale Sensor Networks,” the 1st International Conference on Embedded Networked Sensor Systems (extended abstract in ACM SenSys), November 2003, Los Angeles, California, USA.
16. **Fan Ye**, Gary Zhong, Songwu Lu, Lixia Zhang, “PEAS: A Robust Energy Conserving Protocol for Long-lived Sensor Networks,” the 10th IEEE International Conference on Network Protocols (extended abstract in IEEE ICNP), November 2002, Paris, France, **cited by 301**
17. **Fan Ye**, Alvin Chen, Songwu Lu, Lixia Zhang, “A Scalable Solution to Minimum Cost Forwarding in Large Sensor Networks”, in Proceedings of the 10th International Conference on Computer Communications and Networks (IEEE ICCCN), October 2001, Scottsdale, Arizona, USA, **cited by 111**

Under Submission and Non-refereed

1. **Fan Ye**, Gary Zhong, Songwu Lu, Lixia Zhang, “Energy Efficient Robust Sensing Coverage in Large Sensor Network,” UCLA CS IRL technical report, June 2001, **cited by 65**
2. **Fan Ye**, Songwu Lu, Lixia Zhang, “GRAdient Broadcast: A Robust, Long-lived Sensor Network”, UCLA CS IRL technical report, September 2001, **cited by 40**
3. John Stankovic, Arthur Maccabe, Fan Ye, Philip Levis, “Security,” chapter in the report of NSF workshop on Environmental Cyber-infrastructure Needs for Distributed Sensor Networks, August 2003, San Diego, California, USA.

Patent Applications

3

1. Zhen Liu, Cathy Xia, Hao Yang, Fan Ye, “System and Apparatus for Optimally Trading off the Replication Overhead and Consistency Level in Distributed Applications”, YOR920060281US1, June 2006
2. Zhen Liu, Hao Yang, Fan Ye, “A Method and System for Federated Resource Discovery Service in Distributed Systems”, YOR920060236US1, June, 2006
3. Zhen Liu, Hao Yang, Fan Ye, “Method and System for Resilient Packet Traceback in Wireless Mesh and Sensor Networks”, YOR920060860US1, March 2007

³Authors in alphabetic order for all patents

4. Michael Branson, Fred Douglass, Brad Fawcett, Zhen Liu, William Waller, Fan Ye, "Method for the Interoperation of Virtual Organizations", YOR920070050US1, May 2007
5. Michael Branson, Fred Douglass, Brad Fawcett, Zhen Liu, Anton Riabov, William Waller, Fan Ye, "Method and Apparatus for Cooperative Data Stream Processing", YOR920070049US1, April 2007
6. Michael Branson, Fred Douglass, Brad Fawcett, Zhen Liu, Fan Ye, "Mechanism for Recovery from Site Failure in a Stream Processing System", ROC920070059US1, April 2007
7. Michael Branson, Fred Douglass, Fan Ye, "Mechanism for Execution of Multi-site Jobs in a Data Stream Processing System", ROC920070058US1, May 2007
8. Michael Branson, Fred Douglass, Zhen Liu, Fan Ye, "Method for Inter-site Data Stream Transfer in a Cooperative Data Stream Processing System", ROC920070060US1, May 2007
9. Zhen Liu, Anton Riabov, Fan Ye, "System and Method of Planning for Cooperative Information Processing", YOR920070049US2, May 2007
10. Bo An, Fred Douglass, Brad Fawcett, Anton Riabov, Fan Ye, "Method for Negotiation Management in Collaborative Systems", YOR820070693, Sept 2007
11. Zachary A Garbow, Fred Douglass, Fan Ye, "Mechanism For Encryption Key Management in a Mixed Infrastructure Stream Processing Framework," YOR920070642, Oct 2007

Software

TTDD implementation (in ns2): A system that delivers data to mobile users in large scale sensor networks despite their continuous motion.

GRAB (in Parsec): A robust data delivery service for sensor networks in adverse environment where failures and communication losses are common.

PEAS (in Parsec): An energy management mechanism that allows redundant sensor nodes to sleep to save energy and wake them up when needed.

(available at <http://www.research.ibm.com/people/y/yefan/>)

Computer Skills

Languages and Tools: C/C++, Java, ns2, Parsec, CVS, DDD, Unix shell scripts, Visual C++, Visual Basic, Latex, FORTRAN, MS Office

Operating Systems: Linux, FreeBSD, Solaris/SunOS, Windows 95/98/2000/NT, TinyOS

Professional Activities

Program Committee Member: RTSS Sensor Track 2004, IEEE WASA (International Conference on Wireless Algorithms, Systems and Applications) 2006, IEEE Percom Work-in-Progress Session 2007, IEEE WASA 2007, IEEE Mobiquitous 2007, ISDPE 2007, IEEE Infocom 2008, IEEE ICDCS 2008, IEEE SUTC 2008, MODUS 2008, SSN 2008

Reviewer for journals: ACM Transaction on Networking (ToN), IEEE Transaction on Mobile Computing (TMC), ACM Mobile Networks and Applications Journal (MONET), ACM Wireless Networks Journal (WINET), IEEE Transactions on Vehicular Technology, ACM Transactions on Sensor Networks (TOSN), International Journal on Pervasive Computing and communications (JPCC), IEEE Journal on Communications Surveys and Tutorials, IEEE Transactions on Parallel and Distributed Systems (TPDS), Computer Networks Journal, Journal of Parallel and Distributed Computing (JPDC), EURASIP Journal on Wireless Communications and Networking, EURASIP Journal on Applied Signal Processing

Reviewer for conferences: IEEE INFOCOM, IEEE ICDCS, USENIX OSDI, ACM MobiCom, IEEE WCNC, IEEE Mobiquitous, IEEE Globecom, RTSS, IEEE Broadnets, ICPP, IEEE ICC, ISDPE, IEEE SUTC, MODUS

Attend SensIT PI meeting, Waltham, MA, April, 2000; SensIT PI meeting, St. Petersburg, FL, April, 2001; SIGCOMM 2000, San Diego, CA, August, 2001; MobiCom 2002, Atlanta, Georgia, September, 2002; NSF Environmental Cyber-Infrastructure Workshop on

Sensor Networks, San Diego, August, 2003; MobiCom 2003, San Diego, CA, September, 2003; SenSys 2003, Los Angeles, CA, November, 2003.

Presentations in ICCCN 2001, Scottsdale, AZ, October, 2001; ICNP 2002, Paris, France, November, 2002; SensIT PI meeting, Waltham, MA, November, 2002; IPSN 2003, Palo Alto, CA, April, 2003; ICDCS 2003, Providence, RI, May, 2003; ICDCS 2007, Toronto, Canada, June, 2007. Middleware 2007, Newport Beach, CA, Nov, 2007.