

# Huayong Wang, Ph.D.

---

Singapore-MIT Alliance for Research and Technology (SMART)

Email: huayongw@smart.mit.edu, huayongw@gmail.com

---

## Technical Interests

Non-traditional database (e.g. large scale distributed stream database), cloud computing (programming model, reliability and storage), distributed computing, system software (simulator, compiler and runtime libraries), computer architecture (e.g. transactional memory and many-core systems), parallel architectures and networking in general.

---

## Education

- **Ph.D.**, Computer Science and Technology, Tsinghua University, Beijing, P.R.China, 2006.  
Dissertation: Research on the Dynamic Voltage Scaling Algorithm for Embedded Real-time systems (Best doctoral dissertation award).
  - **M.S.**, Computer Science, Northeastern University, Shenyang, P.R.China, 2002.
  - **B.S.**, Computer Science, Northeastern University, Shenyang, P.R.China, 2000.
- 

## Training

- Micro MBA, IBM Research, Beijing, Aug. 2009.
  - Dale Carnegie Training, IBM Research, Beijing, Mar. ~ Jun. 2009.
  - High Performance Leadership, IBM Research, Beijing, Dec. 2008.
  - Leader Readiness Learning, IBM Research, Beijing, Aug. 2007.
- 

## Professional Experience

- Singapore-MIT Alliance for Research and Technology (<http://smart.mit.edu>), Cambridge (USA) and Singapore, 2010/06/25 ~ Present
  - ✓ Postdoctoral Associate, distributed stream database for intelligent transportation applications.
    - ✧ Designed and implemented MobiStreams – a distributed stream processing system running on both clusters and mobile devices.
    - ✧ Proposed, implemented and evaluated Meteor Shower – a novel fault-tolerance scheme for MobiStreams.
- IBM Research, Beijing, 2006/05/15 ~ 2010/06/04.
  - ✓ Manager, X10 project (<http://x10-lang.org>) – a new programming language for distributed systems with high productivity, 2010/01 ~ 2010/06.
    - ✧ Ported X10 compiler to a new processor with hybrid hardware architecture.
    - ✧ Developed X10 applications and demonstrated the productivity and performance of X10 language.

- ✓ Full-time Staff Research Member and Technical Leader, InfoSphere Streams (<http://www-01.ibm.com/software/data/infosphere/streams>) – a distributed stream computing system which processes large volume of data in real time, 2008/01 ~ 2009/12.
  - ✧ Participated in the design and implementation of a new compiler for stream processing systems.
  - ✧ Designed and implemented the transparent vectorization function in the compiler, which increases the application performance by 3~4 times.
  - ✧ Proposed and evaluated a hardware acceleration approach for stream applications, which increases the application performance nearly 10 times.
  - ✧ Applied the stream computing technology in real-world applications, e.g. intelligent transportation, spam short message filtering, etc.
- ✓ Full-time Staff Research Member and Technical Leader, transactional memory – an innovative hardware architecture for parallel computing, 2007/01 ~ 2007/12.
  - ✧ Proposed a transactional memory system for IBM POWER/PowerPC processors with good balance between cost and performance.
  - ✧ Modeled, implemented and evaluated the transactional memory system on cycle-accurate simulators.
- ✓ Full-time Research Engineer, parallel simulation for multi-core systems, 2006/05 ~ 2006/12.
  - ✧ Parallelized Mambo cycle-accurate simulator through optimistic speculative and conservative approaches.
  - ✧ Participated in IBM Cloud Simulation project. With small accuracy loss (<5%), Cloud Simulation accelerated cycle accurate simulators almost linearly on large scale parallel computing platforms.
- Research Assistant, Tsinghua University, Beijing, P.R.China, 2002 ~ 2005.
  - ✓ Developed dynamic voltage scaling (DVS) algorithms in real-time scheduling systems.
  - ✓ Developed the module in SKYEYE simulator for power consumption estimation.
  - ✓ Taught courses on embedded systems and Linux kernel development.
- Research Assistant, National Engineering Research Center for Computer Software, Northeastern University, P.R.China, 1998 ~ 2002.
  - ✓ Participated in the development of a multi-language web search engine. Implemented a distributed load-balanced crawler and improved the Chinese segmentation algorithm by 1.5 times.
  - ✓ Participated in the development of NetEye firewall 3.0. Implemented a new infrastructure to flexibly manage firewall rules.

---

## Honors and Awards

- IBM Invention Master, 2010.
- Outstanding research accomplishment – System S: High Speed Adaptive Stream Processing & Analytics. IBM Research, 2009.
- Invention achievement award in appreciation and recognition of creative contributions, IBM Research, 2008.
- IBM china outstanding university ambassador in recognition of the outstanding contribution for China university partnership program, 2008.
- Rated among the top 10% contributors in consecutive three years, IBM Research, 2006 ~ 2008.

- Best doctoral dissertation award, Tsinghua University, 2006.
- 

## Publications

- 1) **H. Wang**, L.S. Peh, E. Koukoumidis, T. Shao and M.C. Chan. Meteor Shower: a reliable stream processing system for commodity data centers. Submitted to *38th Int. Conf. Very Large Data Bases (VLDB2012)*, Istanbul, Turkey, Aug. 27-31, 2012.
  - 2) **H. Wang**, F. Calabrese, G. Di Lorenzo and C. Ratti. Transportation mode inference from anonymized and aggregated mobile phone call detail records. In *Proc. 13th Int. IEEE Conf. Intelligent Transportation Syst. (ITSC 2010)*, Madeira Island, Portugal, Sep. 19-22, 2010.
  - 3) Q. Zou, **H. Wang**, R. Soule, M. Hirzel, H. Andrade, B. Gedik and K.L. Wu. From a stream of relational queries to distributed stream processing. In *Proc. 36th Int. Conf. Very Large Data Bases (VLDB 2010)*, Singapore, Sep. 13-17, 2010, pp. 1394-1405.
  - 4) **H. Wang**, Y. Ge, Y. Wang and Y. Zou. Productivity and performance: improving consumability of hardware transactional memory through a real-world case study. In *Proc of Euro-Par*, Ischia-Naples, Italy, Aug. 31-Sep. 3, 2010, pp. 163-174 (LNCS 6272).
  - 5) Z. Liu, C. Wang, Q. Z, and **H. Wang**. Clustering coefficient queries on massive dynamic social networks. In *Proc. 11th Int. Conf. Web-Age Inform. Manage. (WAIM2010)*, Jiuzhaigou, China, Jul 15-17, 2010, pp. 115-126 (LNCS 6184).
  - 6) **H. Wang**, H. Andrade, B. Gedik, and K.L. Wu. A code generation approach for auto-vectorization in the SPADE compiler. In *Proc. 22nd Int. Workshop Languages and Compilers Parallel Computing (LCPC'09)*, Newark, DE, Oct. 8-10, 2009, pp. 383-390 (LNCS 5898).
  - 7) **H. Wang**, H. Andrade, B. Gedik, and K.L. Wu. Auto-vectorization through code generation for stream processing applications. In *Proc. 23rd Int. Conf. Supercomputing (ICS'09)*, Yorktown Heights, NY, Jun. 8-12, 2009, pp. 495-496. (Poster)
  - 8) K. Wang, Y. Zhang, **H. Wang** and X. Shen. Parallelization of IBM mambo system simulator in functional modes. *ACM SIGOPS Operating Syst. Review*, 42(1):71-76, 2008.
  - 9) **H. Wang**, R. Hou, and K. Wang. Hardware transactional memory system for parallel programming. In *Proc. 13th Asia-Pacific Comput. Syst. Architecture Conf. (ACSAC'08)*, Hsinchu, Taiwan, Aug. 4-6, 2008, pp. 1-7.
  - 10) P. Wu, M.M. Michael, C. Praun, T. Nakaike, R. Bordawekar, H.W. Cain, G. Cascaval, S. Chatterjee, S. Chiras, R. Hou, M.F. Mergen, X. Shen, M.F. Spear, **H. Wang**, and K. Wang. Compiler and runtime techniques for software transactional memory optimization. *Concurrency and Computation: Practice and Experience*, 21(1):7-23, 2008.
  - 11) **H. Wang**, Y. Chen and Y. Dai. A soft real-time web news classification system with double control loops. In *Proc. 6th Int. Conf. Web-Age Inform. Manage. (WAIM2005)*, Hanzhou, China, Oct. 11-13, 2005, pp. 81-90 (LNCS 3739).
  - 12) J. Yang, Y. Chen, **H. Wang**, and B. Wang. A Linux kernel with fixed interrupt latency for embedded real-time system. In *Proc. 2nd Int. Conf. Embedded Software and Syst. (ICCESS05)*, Xi'an, China, Dec. 16-18, 2005, pp. 127-134.
  - 13) **H. Wang**, Y. Chen and Y. Dai. An optimal DVS algorithm for real-time system with double scalable processor voltages. *J. Tsinghua University (Natural Science)*, 45(10), 2005.
  - 14) **H. Wang** and Y. Dai. Real-time scheduling scheme for tasks with uncertain execution time. *J. Tsinghua University (Natural Science)*, 45(7): 955-958, 2005.
  - 15) **H. Wang** and Y. Dai. Apply feedback control theory to design soft real-time search engine system. In *Proc. 5th Int. Conf. Web-Age Inform. Manage. (WAIM2004)*, Dalian, China, Jul. 15-17, 2004, pp. 268-279 (LNCS 3129).
  - 16) S. Kang, **H. Wang**, Y. Chen, X. Wang and Y. Dai. Skyeye: an instruction simulator with energy awareness. In *Proc. 1st Int. Conf. Embedded Software and Syst. (ICCESS04)*, Hangzhou, China, 2004, pp. 456-461 (LNCS 3605).
  - 17) **H. Wang** and Y. Dai. An information search interface with soft real-time guarantee. In *Proc. Int. Conf. Inform. Technology: Coding and Computing (ITCC04)*, Las Vegas, NV, Apr. 5-7, 2004, pp. 820-824.
  - 18) X. Zhang, Y. Li, **H. Wang**, and H. Zhao. An improved KNN algorithm applied term feature combination technology for Chinese textual classification. *J. Northeastern University (Natural Science)*, 24(3):229-232, 2003.
  - 19) X. Zhang, **H. Wang**, Y. Li, G. Chang, and H. Zhao. Dynamic vector-space model for Internet textual information categorization. In *Proc. IEEE Int. Conf. Syst., Man, and Cybernetics*, Yasmine Hammamet, Tunisia, Oct. 6-9, 2002, pp. 449-454.
  - 20) X. Zhang, **H. Wang**, G. Chang, and H. Zhao. An autonomous system-based distribution system for web search. In *Proc. IEEE Int. Conf. Syst., Man, and Cybernetics*, Tucson, AZ, Oct. 7-10, 2001, pp. 435-440.
-

## Patents

- 1) **H. Wang**, K. Wang, H. Young. System and method for simulating a multiprocessor system. US Patent Pub. App. No. 20080208558. China Patent No. 200710084322.9.
  - 2) **H. Wang**, K. Wang, H. Young. Apparatus and method for executing rapid memory management unit emulation and full-system simulator. US Patent Pub. App. No. 20080222384. China Patent No. 200710005260.8.
  - 3) **H. Wang**, K. Wang, H. Young. Full-system ISA emulating system and process recognition method. Patent Pub. App. No. 20080270740. China Patent No. 200710104743.3.
  - 4) B. Feng, R. Yan, K. Wang, **H. Wang**. Method and system for analyzing parallelism of program code. US Patent Pub. App. No. 20090031290. China Patent No. 200710109089.5.
  - 5) A. Gheith, **H. Wang**, K. Wang, Y. Zhang. Method, apparatus, and full-system simulator for speeding MMU simulation. US Patent Pub. App. No. 20090119089. China Patent No. 200710167039.2.
  - 6) R. Hou, X. Shen, **H. Wang**. Method and apparatus for implementing transaction memory. US Patent Pub. App. No. 20090119667. China Patent No. 200710169244.2.
  - 7) X. Shen, **H. Wang**, K. Wang. Method and system for handling transaction buffer overflow in a multiprocessor system. US Patent Pub. App. No. 20090144524. China Patent No. 200710196184.3.
  - 8) W. Shen, P. Shao, Y. Li, X. Chang, Y. Ge, **H. Wang**, H. Zou. Method of processing instructions in pipeline-based processor and corresponding processor. US Patent Pub. App. No. 20090193424. China Patent No. 200810003557.5.
  - 9) X. Shen, **H. Wang**, W. Shen, P. Shao. Method and device for performing copy-on-write in a processor. US Patent Pub. App. No. 20090248984. China Patent No. 200810086951.X.
  - 10) **H. Wang**, C.B. Hall, Y. Wang, Z. Liang, X. Shen. System enabling transactional memory and prediction-based transaction execution method. US Patent Pub. App. No. 20090292884. China Patent No. 200810109178.4.
  - 11) R. Hou, Z.Y. Liu, **H. Wang**, Y.Q. Wang. Method and system for loading status control of DLL. US Patent Pub. App. No. 20100106950. China Patent No. 200810170786.6.
  - 12) H.W. Cain, R. Hou, X. Shen, **H. Wang**. Method and system for a sharing buffer. US Patent Pub. App. No. 20100138571. China Patent No. 200810181608.3.
  - 13) X.T. Chang, **H. Wang**, K. Wang, Y. Zhang. Simulator and simulating method for running guest program in host. US Patent Pub. App. No. 20100161875. China Patent No. 200810185788.2.
  - 14) **H. Wang**. Method and apparatus for implementing a transactional store system using a helper thread. US Patent Pub. App. No. 20100186015. China Patent No. 200910005955.5.
  - 15) Y. Ge, R. Hou, **H. Wang**. Fast context save in transactional memory. US Patent Pub. App. No. 20100217945. China Patent No. 200910008371.3.
  - 16) R. Hou, Z. Liu, **H. Wang**, Y. Wang, Q. Zou, Y. Zou. Method and system for sampling input data. US Patent Pub. App. No. 20100281310. China Patent No. 200910136950.6.
  - 17) H. Andrade, B. Gedik, R. Hou, **H. Wang**, K. Wu. Incrementally constructing executable code for component-based applications. US Patent Pub. App. No. 20100293533.
  - 18) H. Andrade, B. Gedik, **H. Wang**, K. Wu. Use of vectorization instruction sets. US Patent Pub. App. No. 20100293534.
  - 19) X. Shen, **H. Wang**, R. Hou, Y. Li. A cache with novel replacement algorithm for transactional memory system. China Patent No. 200810083947.8.
  - 20) P. Shao, X. Shen, R. Hou, **H. Wang**, Y. Ge. Method, apparatus and system for sharing cache dynamically in multi-core processors. China Patent No. 200810083946.3.
  - 21) K. Wang, **H. Wang**, Y. Zhang, X. Shen. Prime-based ordering method for transactional memory. China Patent No. 200810109302.7.
  - 22) **H. Wang**, W. Shen, R. Hou, X. Shen. Transactional memory system and control method thereof. China Patent No. 200810127541.5.
  - 23) **H. Wang**, R. Hou, K. Wang, X. Shen. An apparatus and method for lazy fine-grained copy-on-write. China Patent No. 200810131937.7.
  - 24) Y. Zhu, Y. Zhang, W. Shang, J. Zhou, **H. Wang**, C. Ying. Cell handover based virtual loop detector for transportation systems. China Patent No. 200910247123.4.
- The following patents are filed but not available online yet.
- 25) Y. Zhu, Z. Xiang, **H. Wang**, W. Shang, J. Zhou, C. Ying. System and method of monitoring traffic with mobile network based assistant positive positioning.
  - 26) H. Andrade, B. Gedik, Q. Zou, **H. Wang**, R. Soule, M. Hirzel, K.-L. Wu. Transforming relational queries into stream processing.
  - 27) H. Andrade, B. Gedik, Q. Zou, **H. Wang**, R. Soule, M. Hirzel, K.-L. Wu. Proxing ODBC calls.

---

## Professional Activities

- Co-chair, IBM China Research Lab invention development team (patent review board), Apr. 2008 ~ May, 2010.
  - Core team member, Global Technology Outlook core team in IBM China Research Lab, 2009.
  - IEEE and ACM Membership.
- 

## Presentations and Invited Talks

- 1) “Stream computing in telecom industry”, Invited Talk, China Mobile Research Institute, April 1, 2009
- 2) “Stream computing for banking industry”, Invited Talk, China Unionpay Research Lab, December 24, 2008
- 3) “What is happening in VLSI industry”, Invited Tutorial, Beijing Univ. of Posts and Telecommunications, May 14, 2008
- 4) “Transactional memory: from research to product”, Invited Tutorial, Tsinghua Univ., November 28, 2007