
Bingsheng He

Ph.D., HKUST, 2008

Email: hebingsheng(at)msn(dot)com

Tel: (86) 135 5298 6932, (8610)58925922

Address: Microsoft Research Asia 4/F, Beijing Sigma Center No.49, Zhichun Road, Hai Dian District Beijing China 100080

Education

- | | |
|-------------------|--|
| 08/2003 — 07/2008 | Ph. D. of Computer Science
Computer Science and Engineering Department
Hong Kong University of Science & Technology (HKUST)
Supervisor: Dr. Qiong Luo |
| 09/1999 — 08/2003 | B.E. of Computer Science and Engineering
B.B.A. of Business Administration (secondary major)
Shanghai Jiao Tong University, China |
| 09/1993— 06/1999 | Zhongshan Memorial Middle School
(named after Dr. Sun Yat-Sen)
Zhongshan, Guangdong, China |

Working Experiences

Post-Doc Researcher,

Microsoft Research at Asia

09/2008 — present System Research Group

Research or Teaching Assistant,

Hong Kong University of Science & Technology

09/2003 — 08/2008 Under the supervision of Dr. Qiong Luo

Research Interests

- Database Management Systems (particularly architecture-aware database systems)
- Data Management in Data Centers
- GPGPU (General-Purpose Computation on Graphics Processors)
- High Performance Computing
- XML and Graph Databases

Research and Professional Experiences

- **Current research at Microsoft Research at Asia, 09/2008 – present**
My current project is data management in data centers.
- **Ph.D. study at Hong Kong University of Science and Technology, 08/2003—08/2008**
My Ph.D. study focused on architecture-aware databases, including cache-centric techniques and parallel algorithms on CPUs (Central Processing Units) and GPUs (Graphics Processor Units). I was the principal developer in the following projects:
 - Cache-conscious techniques for query processing techniques in XML databases

- (ICDE 2005, TKDE 2006, and TKDE 2007).
- A cache-oblivious in-memory query processor for relational databases (CIKM 2006, CIDR 2007, TODS 2008).
 - An in-memory relational query co-processing using GPUs (SuperComputing 2006, SIGMOD 2008).
 - Efficient and parallel access methods for relational databases on multi-core processors (in submission).
 - Mars: a MapReduce framework on the GPU (PACT 2008).
 - **Summer internship at Microsoft Research at Asia**, summer 2007
 - Analyzed and modeled the throughput of the index file manager of the Live search.
 - Improved the throughput of the index file manager using the emerging hardware.
 - My mentor is Dr. Zhang Zheng.
 - **UG (Undergraduate) thesis project at Shanghai Jiao Tong University**, 1999
The title of my UG thesis project was “the performance analysis of the stock market network”. The tutors for my UG thesis project were Mr. Zhaodong Wang (Chief Architect of HP Corporation) and Prof. Yong Yu.
 - **Summer practice in American Standard SRC at Shanghai**, 1999
 - Developed programs to manipulate the data from the backend database.

Publications¹

Summary

My publications appear mainly in two kinds of forums. Here is a summary of the major forum and the number of papers in the forum:

- Databases: TODS (1), TKDE (2), SIGMOD (1), SIGMOD Demo (2), ICDE(2), CIDR (1).
- System: SuperComputing (1), PACT (1), PPOPP (1).

Journal Articles

[DB*] **Bingsheng He** and Qiong Luo, Cache-Oblivious Databases: Limitations and Opportunities. *TODS 2008: ACM Trans. Database Syst.* 33, 2 (Jun. 2008), 1-42.

[DB*] **Bingsheng He**, Qiong Luo and Byron Choi. Adaptive Index Utilization in Memory-Resident Structural Joins. *TKDE 2007: IEEE Transactions on Knowledge and Data Engineering*, vol. 19, no. 6, pp. 772-788, June, 2007.

[DB*] **Bingsheng He**, Qiong Luo and Byron Choi. Cache-Conscious Automata for XML Filtering. *TKDE 2006: IEEE Transactions on Knowledge and Data Engineering*, vol. 18, no. 12, pp. 1629-1644, Dec., 2006.

Conference and Workshop Papers

[SYSTEM*] Ke Yang, **Bingsheng He**, Qiong Luo, Pedro Sander, Jiaoying Shi. Stack-Based

¹ Notations: [DB] for database forums/tracks; [SYSTEM] for system forums/tracks; * for major forums/tracks.

Parallel Recursion on Graphics Processors. (Poster) To appear in PPOPP09: 14th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming.

[DB*] Yinan Li, **Bingsheng He**, Qiong Luo and Ke Yi. Tree Indexing on Flash Disks. (short paper) To appear in IEEE International Conference on Data Engineering (ICDE), 2009.

[SYSTEM*] **Bingsheng He**, Wenbin Fang, Qiong Luo, Naga K. Govindaraju, and Tuyong Wang. Mars: A MapReduce Framework on Graphics Processors. *PACT08: IEEE International Conference on Parallel Architecture and Compilation Techniques 2008*.

[DB] Hongyu Guo, **Bingsheng He**, Yifan He, Qiong Luo, Bo Peng, and Xiangye Xiao. Frequent Pattern Mining on Graphics Processors. *ACM SIGMOD 2008 undergraduate poster competition*.

[DB*] **Bingsheng He**, Ke Yang, Rui Fang, Mian Lu, Naga K. Govindaraju, Qiong Luo, and Pedro V. Sander. Relational Joins on Graphics Processors. *SIGMOD 2008: ACM SIGMOD 2008*.

[SYSTEM*] **Bingsheng He**, Naga K. Govindaraju, Qiong Luo, Burton Smith. Efficient Gather and Scatter Operations on Graphics Processors. *SC2007: ACM/IEEE SuperComputing 2007*.

[DB/SYSTEM*] **Bingsheng He** and Qiong Luo. Cache-Oblivious Query Processing. *CIDR 2007: The 3rd Biennial Conference on Innovative Data Systems Research*.

[DB] **Bingsheng He** and Qiong Luo. Cache-Oblivious Nested-Loop Joins. *CIKM 2006: ACM Fifteenth Conference on Information and Knowledge Management*.

[DB] Zi Lin, **Bingsheng He**, and Byron Choi. A Quantitative Summary of XML Structures. *ER 2006: 25th International Conference on Conceptual Modeling*.

[DB*] **Bingsheng He**, Qiong Luo and Byron Choi. Cache-Conscious Automata for XML Filtering. *ICDE 2005: 21st International Conference on Data Engineering*.

[DB] **Bingsheng He**, Yinan Li, Qiong Luo and Dongqing Yang. A General Framework for Improving Query Processing Performance on Multi-Level Memory Hierarchies. *ACM SIGMOD Third International Workshop on Data Management on New Hardware 2007*.

[DB] Ke Yang, **Bingsheng He**, Rui Fang, Mian Lu, Naga K. Govindaraju, Qiong Luo, Pedro Sander, Jiaoying Shi. In-Memory Grid Files on Graphics Processors. *ACM SIGMOD Third International Workshop on Data Management on New Hardware 2007*.

[DB] Qiong Luo, Lionel M. Ni, **Bingsheng He**, Hejun Wu, and Wenwei Xue. MEADOWS: Modeling, Emulation, and Analysis of Data of Wireless Sensor Networks. *International Workshop on Data Management for Sensor Networks (DMSN'04) in conjunction with VLDB 2004*.

[SYSTEM] Wenwei Xue, **Bingsheng He**, Hejun Wu, Qiong Luo. The HKUST Frog Pond - A Case Study of Sensory Data Analysis. *NPC 2004: IFIP International Conference on Network and Parallel Computing*.

[SYSTEM] Hejun Wu, Qiong Luo, Pei Zheng, **Bingsheng He**, and Lionel M. Ni. Accurate Emulation of Wireless Sensor Networks. *NPC 2004: IFIP International Conference on Network and Parallel Computing*.

System Demonstration

[DB*] **Bingsheng He**, Yinan Li, Qiong Luo, Dongqing Yang. EaseDB: A Cache-Oblivious In-Memory Query Processor. *ACM SIGMOD 2007*.

[DB*] Rui Fang, **Bingsheng He**, Mian Lu, Ke Yang, Naga K. Govindaraju, Qiong Luo, Pedro Sander. GPUQP: Query Co-Processing using Graphics Processors. *ACM SIGMOD 2007*.

Technical Reports

[DB*] **Bingsheng He et al**, a paper on query coprocessing. *Invited submission to a top journal (details omitted due to anonymity)*.

[DB] Qiong Luo, Hejun Wu, Wenwei Xue, and **Bingsheng He**. Benchmarking In-network Sensor Query Processing. *Technical Report HKUST-CS05-09, Department of Computer Science, HKUST, June 2005*.

[DB] **Bingsheng He** and Qiong Luo. Cache-Oblivious Hash Joins. *Technical report HKUST-CS06-04, Department of Computer Science and Engineering, HKUST, July 2006*.

Book Chapter

Peng Ding, Fang Liu, Zhifeng Shao, **Bingsheng He** (Ed.). STRUCTS Technology and Its Applications in Web (Chinese). ISBN 7-302-08105-0. Tsinghua University Press, March 2004.

Professional Activities

- PC member for DASFAA 2009 (demo), WCMT 2009 (with WAIM/APWeb 2009).
- External reviewer for NDBC 2004, ICDE 2005, ICDE 2009, VLDB 2004, VLDB 2006, TKDE 08, TOIS 08.
- Student helper for ICDM2006, Hong Kong.
- Seminar coordinator at the system group of MSRA.

Honors and Awards

- IBM PhD Fellowship 2007/08.
(News: http://www.cse.ust.hk/News/IBM_Fellowship2007/).
- Graduate Student Scholarship of CIDR 2007 (The 3rd Biennial Conference on Innovative Data Systems Research).
- Outstanding undergraduate student of Shanghai City, 2003.
- First class scholarship of Shanghai Jiao Tong University, 1999—2003.
- First class studentship of Bank of Communications, China, at Shanghai Jiao Tong University, 1999.

References

Available upon request