

Bin Shao

CONTACT INFORMATION Microsoft Research Asia Office: +86 (10) 59174329
Bldg 2, No. 5 Dan Ling Street, Beijing, PRC 100080 E-mail: binshao@microsoft.com

SUMMARY Bin Shao is a lead researcher at Microsoft Research Asia (Beijing, China). He joined Microsoft after receiving his Ph.D. degree from Fudan University in July 2010. Bin Shao is leading the Microsoft Graph Engine project, which is a distributed in-memory large graph processing engine. His research interests include machine learning techniques, in-memory databases, distributed systems, and parallel graph processing. The results of his research have appeared in top conferences and leading journals.

EDUCATION School of Computer Science, Fudan University, Shanghai, China
Ph.D. Sep. 2005 – Jul. 2010
School of Computer Science and Technology, Shandong University, Jinan, China
B.E. Sep. 2001 – Jul. 2005

HONOURS AND AWARDS Gold Star Award, Microsoft, 2011
Outstanding Graduates Awards, Shanghai, 2010
OOCL(Orient Overseas Container Line) Scholarship, Fudan University, 2009
Scholarship for Selected First-Year Graduate Students, Fudan University, 2005
Excellent Graduation Thesis, Shandong University, 2005
Outstanding Graduates Awards, Shandong University and Shandong Province, 2005
First Prize, Shandong College Student Software Design Contest, 2004
Scholarship (2nd Prize or Better) Every Year, Shandong University, 2002-2004
CVIC(China Venturetech Investment Corporation) Software Engineering Scholarship, 2002

SELECTED PUBLICATIONS *Bin Shao*, Haixun Wang, and Yatao Li. Trinity: A Distributed Graph Engine on a Memory Cloud. In SIGMOD 2013: Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data, New York, New York, USA, June 22-27, 2013, pages 505-516.

Hongbin Ma, *Bin Shao*, Yanghua Xiao, Jeff Liang Chen, and Haixun Wang. G-SQL: Fast Query Processing via Graph Exploration. In PVLDB 2016: Proceedings of the 42th International Conference on Very Large Data Bases, New Delhi, India, September 5-9, 2016, pages 900-911.

Bin Shao, Haixun Wang and Yanghua Xiao. Managing and Mining Large Graphs: Systems and Implementations (Tutorial). In SIGMOD 2012: Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data, Scottsdale, Arizona, USA, May 20-24, 2012, pages 589-592.

Zichao Qi, Yanghua Xiao, *Bin Shao*, Haixun Wang. Toward a Distance Oracle for Billion-Node Graphs. In PVLDB 2014: Proceedings of the 40th International Conference on Very Large Data Bases, Hangzhou, China, September 1-5, 2014, pages 61-72.

Lu Wang, Yanghua Xiao, *Bin Shao*, Haixun Wang. How to Partition a Billion-Node Graph. In ICDE 2014: Proceedings of the 30th IEEE International Conference on Data Engineering, Chicago, IL, USA, March 31-April 4, 2014, pages 568-579.

Huanhuan Xia, Tun Lu, *Bin Shao*, Guo Li, Xianghua Ding, Ning Gu. A partial Replication Approach for Anywhere Anytime Mobile Commenting. In CSCW 2014: Proceedings of the 17th ACM conference on CSCW and social computing, Baltimore, MD, USA, February 15-19, pages 530-541.

Kai Zeng, Jiacheng Yang, Haixun Wang, *Bin Shao*, and Zhongyuan Wang. A Distributed Graph Engine for Web Scale RDF Data. In PVLDB 2013: Proceedings of the 39th international conference on Very Large Data Bases, Riva del Garda, Trento, August 26-30, pages 265-276.

- Zhao Sun, Hongzhi Wang, Haixun Wang, *Bin Shao*, and Jianzhong Li. Efficient Subgraph Matching on Billion Node Graphs. In PVLDB 2012: Proceedings of the 38th international conference on Very Large Data Bases. 5(9):788-799, 2012.
- Huanhuan Xia, Tun Lu, *Bin Shao*, Xianghua Ding, Ning Gu. Hermes: On Collaboration across Heterogeneous Collaborative Editing Services in the Cloud. In CSCWD 2014: 18th IEEE International Conference on CSCW in Design, Hsinchu, Taiwan, May 21-23, 2014, pages 655-660.
- Liang He, *Bin Shao*, Yatao Li, Enhong Chen. Distributed Real-Time Knowledge Graph Serving. In BigComp 2015: Second International Conference on Big Data and Smart Computing, Jeju island, Korea, February 9-12, 2015. Invited Paper.
- Ho Lee, *Bin Shao*, U Kang. Fast graph mining with HBase. Inf. Sci. 315, C, 2015, pages 56-66.
- Yatao Li, *Bin Shao*. Hash storage based large graph generator. Communications of the CCF (Chinese). 8(11): 16 -20, 2012.
- Bin Shao*, Du Li, Tun Lu and Ning Gu. An Operational Transformation Based Synchronization Protocol for Web 2.0 Applications. In CSCW 2011: Proceedings of the 2011 ACM Conference on Computer Supported Cooperative Work, Hangzhou, China, Mar. 19-23, 2011.
- Bin Shao*, Du Li, and Ning Gu. An Algorithm for Selective Undo of Any Operation in Collaborative Applications. In Group 2010: Proceedings of Group 2010 Conference, Sanibel Island, FL., USA, Nov. 7-10, 2010, pages 131-140.
- Bin Shao*, Tun Lu, and Ning Gu. Key Techniques of Consistency Maintenance in Real-time Collaboration. Computer Engineering, vol.36, no.22, pages 1-6, Nov. 2010. Invited Paper.
- Bin Shao*, Du Li, and Ning Gu. A Sequence Transformation Algorithm for Supporting Cooperative Work on Mobile Devices. In CSCW 2010: Proceedings of the 2010 ACM Conference on Computer Supported Cooperative Work, Savannah, GA, USA, Feb. 6-10, 2010, pages 159-168.
- Bin Shao*, Du Li, and Ning Gu. A Fast Operational Transformation Algorithm for Mobile and Asynchronous Collaboration. IEEE Transactions on Parallel and Distributed Systems, vol.21, no.12, pages 1707-1720, Dec. 2010.
- Bin Shao*, Du Li, and Ning Gu. ABTS: A Transformation-based Consistency Control Algorithm for Wide-area Collaborative Applications. In CollaborateCom 2009: The 5th International Conference on Collaborative Computing: Networking, Applications and Worksharing, Washington D.C., USA, Nov. 11-14, 2009. Invited Paper.
- Bin Shao*, Du Li, and Ning Gu. An Optimized String Transformation Algorithm for Real-time Group Editors. In ICPADS 2009: The 15th International Conference on Parallel and Distributed Systems, Shenzhen, China, Dec. 8-11, 2009, pages 376-383.
- Liping Gao, *Bin Shao*, Lin Zhu, Tun Lu, and Ning Gu. Maintaining Time and Space Consistencies in Hybrid CAD Environments: Framework and algorithms. Computers in Industry, Vol. 59, Issue 9, Dec. 2008, pages 894-904.
- Sili Zhao, Du Li, Hansu Gu, *Bin Shao*, and Ning Gu. An Approach to Sharing Legacy TV/Arcade Games for Real-time Collaboration. In ICDCS 2009: The 29th International Conference on Distributed Computing Systems, Montreal, Canada, Jun. 22-26, 2009, pages 165-172.
- Liping Gao, *Bin Shao*, Tun Lu, and Ning Gu. Maintaining Semantic Intention of Step-wise Operations in Replicated CAD Environments. In CSCWD 2008: The 12th International Conference on CSCW in Design, Xi'an, China, Apr. 16-18, 2008, pages 154-159.

Liping Gao, *Bin Shao*, and Ning Gu. Separating Data and View: Support View-wandering Between Different Trades During Engineering Design. In CSCWD 2007: The 11th International Conference on CSCW in Design, Melbourne, Australia, Apr. 26-28, 2007, pages 36-41.

Hansu Gu, Qiwei Zhang, and *Bin Shao*. Making AutoCAD Collaborative : Implementation and Application of CoAutoCAD. In ICPCA 2007: Proceedings of 2nd International Conference on Pervasive Computing and Applications, Birmingham, England, Jul. 26-27, 2007, pages 168-173.

PATENT
APPLICATIONS

Bin Shao, Du Li, Ning Gu. China Patent Application 201010176176.4, A Selective Undo Method based on Operation Effects Relation Order. Publication date: Oct. 6, 2010.

Bin Shao, Haixun Wang, Wei Fang. US Patent Application 2015/0120775 A1, Answering relational database queries using graph exploration. Application Number: 14/062036. Publication date: Apr. 30, 2015.