

EDUCATION

- Ph.D., Computer Science, Purdue University, West Lafayette, IN**
Emphasis: Multimedia/Streaming data management systems
Date: August 2007
- Certificate of Applied Management Principles (mini-MBA), Krannert School of Management, Purdue University, West Lafayette, IN**
Date: May 2004
- M.S., Computer Science, Purdue University, West Lafayette, IN**
Date: May 2003
- M.S., Horticulture, Purdue University, West Lafayette, IN**
Date: December 2000
- B.S., Horticulture, Beijing Agricultural University, Beijing, China**
Date: July 1993

EMPLOYMENT

Associate Professor, Department of Computer Science and Engineering, University of South Florida, 08/2013 - date

Assistant Professor 08/2007 – 07/2013

- Published and presented research results in multiple top international journals and conferences
- Supervised undergraduate/Master/PhD level students for their theses/dissertations
- Taught undergraduate/graduate level courses in database systems
- Projects involved:
 1. *Database-Centric Molecular Simulation*
 2. *Power-aware database management systems*
 3. *Self-tuning databases based on feedback control theory*
 4. *Building a CUDA Teaching/Research Center at USF College of Engineering*
 5. *Push-based data management systems on GPUs*
 6. *Acquisition, maintenance of the GAIVI computer cluster*

Research Assistant, CS Department, Purdue University, 5/2001- 7/2007

- Published and presented research results in multiple top international journals and conferences
- Supervised undergraduate/Master level students for their independent research/study projects
- Projects involved:
 1. *Quality of Service Control in Emerging Data Management Systems*
 2. *Knowledge Projection - A Dynamic Maintenance System for Navy Warships*
 3. *VDBMS: A Video DBMS with Content-based Retrieval Capabilities*
 4. *Cost-effective Media Streaming: A Peer-to-peer Approach*

Summer Intern, R&D division, Micro Data Bases, Inc., West Lafayette, IN, 5/2001 - 8/2001

- Key member of the team for developing *Fiji* – a large-object supportive framework in DBMS
- Intensive DBMS kernel development and performance evaluation.

Co-founder and Manager, Hunan Wilson Consultants, Ltd., Changsha, China, 12/1993 – 7/1997

- Founded and managed a joint venture in the consulting business
- Explored new business model and open branches in 3 cities in China

GRANTS

- Co-Investigator, *Research Education in Biostatistics: A Life Changing Exposure to Data Science via Epidemiology, Clinical Trials, and Public Health*, National Heart, Lung, and Blood Institute, NIH, April 2016 – March 2019, \$700,000.

Yicheng Tu

- PI, *II-New: A research platform for heterogeneous, massively parallel computing*, \$679,798, National Science Foundation (NSF), 07/01/2015 - 06/30/2018.
- PI, *CUDA Research Center at the University of South Florida*, \$8,000 equipment donation. NVIDIA Corp., NVIDIA Research Center Program, October 2013 – December 2015.
- PI, *CAREER:Enabling High-Throughput Data Management for Scientific Domains*, \$499,883, National Science Foundation, June 2013 – May 2018.
- PI, *Building a CUDA Teaching Center at University of South Florida*, \$11,196 (\$6,696 cash + \$4,500 equipment donation) plus \$6,696 USF matching funds, July 2012 – December 2017.
- PI, *Towards Stream Data Management on GPUs*. \$2,500 equipment donation. NVIDIA Corp., NVIDIA Academic Partnership Program, December 2011.
- PI (leading), *III: Small: Collaborative Research: Making Databases Green: An Energy-Aware DBMS Approach*, National Science Foundation, \$461,678 (PI's budget: \$230,850, USF Co-PI: Bo Zeng, Ohio State University PI: Xiaorui Wang), July 2011 – June 2014.
- PI, *Database-Centric Data Analysis of Molecular Simulations (R01)*, National Institutes of Health, \$875,262 (PI's budget: \$547,993, Co-PI: Sagar Pandit), April 2010 – March 2015.
- PI, *Reducing the Energy Consumption of IT: A Focus on Data Centers and End User Devices*, Florida Energy Systems Consortium USF Seed Grant, \$50,000 (PI's budget: \$24,601, Co-PI: Ken Christensen), May 2010 – June 2011
- PI, *Automatic Database Tuning: A Control-Theoretical Approach*, University of South Florida New Researcher Grant, \$6,000, May 2008 – May 2009.

STUDENTS SUPERVISED

PhD students

- Anand Kumar (graduated in December 2013)
- Vladimir Grupcev (graduated in May 2015)
- Purushottam Panta (passed qualifiers in Spring 2012)
- Peyman Behzadnia (started in Fall 2010, passed qualifiers)
- Ran Rui (started in Fall 2012, passed qualifiers)
- Hao Li (started in Fall 2013, passed qualifiers)
- Mehrad Eslami (started in Spring 2014, passed qualifiers)
- Zhila Nouri (started in Fall 2014, passed qualifiers)
- Faisal Qarah (started in Fall 2014, passed qualifiers)
- Chengcheng Mou (started in Fall 2015)
- Napath Pitaksrianan (started in Fall 2015)
- Tao Hou (started in Fall 2015)
- Zichen Xu (started in Fall 2009, now an assistant professor in Nanchang University, China)
- Jin Huang (visiting student from UT Arlington, Fall 2011)

MS students

- *Current*: Jinghan Meng
- *Graduated*: Iliiazbek Akhmedov (December 2016), Sam Cyrus (August 2016), Parneet Kaur (August 2016), Chengcheng Mou (Summer 2015), Ryan Wheeler (Fall 2015), Meryem Berrada (Spring 2015), Nadim Bou Zeidan (Fall 2014), Di Yu (Spring 2013), Tiffany Burrell (Spring 2010), Sadhana Sharma (Fall 2009), Zichen Xu (Summer 2009)

Undergraduate students

- *REU students*: Anthony Casagrande, Steve Shifke, Evan Kroske, Andrew Jones, Miguel Rodriguez, Exequiel Bravo, Micca Osmar, Amarilys Mendez, Andrey Shipalov, Bhargava Kondaveeti, Laura Reider
- *Honor's thesis*: Ivan Dyedov
- *Senior design project*: Gian Framingheddu

REFEREED PUBLICATIONS (underlined names are senior authors)

Refereed Journal Publications

1. Yin Lu, Dan Shen, Maxwell Pietsch, Chetan Nagar, Zayd Fadli, Hong Huang, **Yi-Cheng Tu**, and Feng Cheng. A novel algorithm for analyzing drug-drug interactions from MEDLINE literature. *Scientific Reports* 5, 17357, November 2015. (IF: 5.58)
2. Z. Xu, **Y. Tu**, and X. Wang. Online Energy Estimation of Relational Operations in Database Systems. *IEEE Transactions on Computers (TC)* 64(11):3223-3236,

3. A. Kumar, V. Grupcev, M. Berrada, J. Fogarty, **Y. Tu**, X. Zhu, S. Pandit, and Y. Xia. DCMS: A data analytics and management system for molecular simulation. *Journal of Big Data* 2:9, November 2014.
4. A. Kumar, V. Grupcev, Y. Yuan, **Y. Tu**, Jin Huang, and G. Shen. Computing Spatial Distance Histograms for Large Scientific Datasets On-the-fly. *IEEE Transactions on Knowledge and Data Engineering (TKDE)* 26(10):2410-2424, October 2014.
5. **Y. Tu**, X. Wang, and B. Zeng. A System for Energy-Efficient Data Management. *ACM SIGMOD Record* 43(1):21-26.
6. J. Huang, F. Nie, **H. Huang**, **Y. Tu**, and Y. Lei. Social Trust Prediction Using Heterogeneous Networks. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 7(4):17, November 2013.
7. C. H. Nadungodage, **Y. Xia**, J. Lee, and **Y. Tu**. Hyper-Structure Mining of Frequent Patterns in Uncertain Data Streams. *Knowledge and Information Systems (KAIS)* 37(1):219-244, October 2013.
8. V. Grupcev, Y. Yuan, **Y. Tu**, J. Huang, S. Chen, S. Pandit, and M. Weng. Approximate Algorithms for Computing Spatial Distance Histograms with Accuracy Guarantees. *IEEE Transactions on Knowledge and Data Engineering (TKDE)* 25(9):1982-1996. September 2013.
9. Z. Xu, **Y. Tu**, and X. Wang. PET: Reducing Database Energy Cost via Query Optimization (software demo). *Proceedings of the Very Large Database Endowment (VLDB)*. 5(12):1954-1957, August 2012.
10. S. Chen, **Y. Tu**, and Y. Xia. Performance Analysis of Dual-Tree Algorithms for Computing Spatial Distance Histograms. *Very Large Data Base (VLDB) Journal*. 20(4):471-494, August 2011.
11. **R. Cheng**, B. Kao, S. Prabhakar, A. Kwan, and **Y. Tu**. Filtering Data Streams for Entity-based Continuous Queries. *IEEE Transactions on Knowledge and Data Engineering (TKDE)* 22(2):234-248, February 2010.
12. **H. Fang**, Q. Wang, **Y. Tu**, and M. F. Horstemeyer. An efficient non-dominated sorting method for evolutionary algorithms. *Journal of Evolutionary Computation*. 16(3):355-384, Fall 2008.
13. **Y. Tu**, J. Yan, G. Shen and S. Prabhakar. Multi-Quality Data Replication in Multimedia Databases. *IEEE Transactions on Knowledge and Data Engineering (TKDE)* 19(5):679-694, May 2007.
14. **L. Qu** and **Y. Tu**. Change Point Estimation of Bi-Level Functions. *Journal of Modern Applied Statistical Methods*. 5(2):347-355, November 2006.
15. **Y. Tu**, J. Sun, M. Hefeeda, and S. Prabhakar. An Analytical Study of Peer-to-Peer Media Streaming Systems. *ACM Transactions on Multimedia Computing, Communications, and Applications (TOMCCAP)*. 1(4):354-376, November 2005.
16. W. Aref, A. Catlin, **A. Elmagarmid**, J. Fan, M. Hammad, I. Ilyas, M. Marzouk, S. Prabhakar, **Y. Tu** and X. Zhu. VDBMS: A Testbed Facility for Research in Video Database Benchmarking. *ACM/Springer Multimedia Systems Journal*. 9(6):575-585, June 2004.

Book Chapters

1. **Y. Tu** and Gang Ding. Control-Based Database Tuning Under Dynamic Workloads. Vol. I., *Encyclopedia of Data Warehousing and Mining*, 2nd edition, pp.333-338, 2009.

Refereed Conference Publications

1. Jinghan Meng and **Yi-Cheng Tu**. Flexible and Feasible Support Measures for Mining Frequent Patterns in Large Labeled Graphs. Accepted to *ACM International Conference on Management of Data (SIGMOD)*, Raleigh, NC, USA., May 14-19, 2017.
2. N. Pitaksirianan, Z. Nouri, and **Y. Tu**. Efficient 2-body statistics computation on GPUs: Parallelization and Beyond. In *Procs. 45th International Conference on Parallel Processing (ICPP)*. pp. 380-385, Philadelphia, PA, USA., August 16-19, 2016.
3. P. Behzadnia, W. Yuan, B. Zeng, **Y. Tu**, and X. Wang. Dynamic Power-Aware Disk Storage Management in Database Servers. In *Procs. 27th International Conference on Database and Expert Systems Applications (DEXA)*. pp. 315-325, Porto, Portugal, September 5-8, 2016.
4. A. Kumar, J. Ligatti, and **Y. Tu**. Query Monitoring and Analysis for Database Privacy - A Security Automata Model Approach. In *Procs. of 16th International Conference on Web Information Systems Engineering (WISE)*, pp. 2458-2472, Miami, FL, USA., November 2015.
5. R. Rui, H. Li, and **Y. Tu**. Join algorithms on GPUs: A revisit after seven years. In *Procs. of 3rd IEEE International Conference on Big Data (BigData)*, pp. 2541-2550, Santa Clara, CA, USA., October 2015.
6. V. Grupcev, **Y. Tu**, J. C. Fogarty, S. Pandit. Push-based system for molecular simulation data analysis. In *Procs. of 3rd IEEE International Conference on Big Data (BigData)*, pp. 1775-1784,

- Santa Clara, CA, USA., October 2015.
7. H. Li, D. Yu, A. Kumar, and **Y. Tu**. Performance Modeling in CUDA Streams – A Means for High-Throughput Data Processing. *IEEE International Conference on Big Data*, pp.301-310, October 2014.
 8. A. Kumar, X. Zhu, **Y. Tu**, and S. Pandit. Compression of Molecular Simulation Datasets. *International Conference on Intelligence Science and Big Data Engineering (ISciDE)*, pp.22-29, July 2013.
 9. Miguel Rodriguez, Daladier Jabba, Elias Nino, Carlos Ardila, and Yi-Cheng Tu. Automata Theory Based Approach to the Join Ordering Problem in Relational Database Systems. Procs. of *International Conference on Data Management Technologies and Applications (DATA)*. pp.257-265, Reykjavik, Iceland. July 29-31, 2013. (**Best Paper Nominee**).
 10. **Y. Tu**, A. Kumar, D. Yu, R. Rui, and R. Wheeler. Data Management Systems on GPUs: Promises and Challenges. *25th Scientific and Statistical Database Management Conference (SSDBM)*. Baltimore, Maryland, USA., July 29-31, 2013.
 11. Zichen Xu, Xiaorui Wang, and Yi-Cheng Tu. Power-Aware Throughput Control for Database Management Systems. Procs. of *10th International Conference on Autonomic Computing (ICAC)*. June 2013.
 12. Zichen Xu, Yi-Cheng Tu, and Xiaorui Wang. Dynamic Energy Estimation of Query Plans in Database Systems. Procs. of *33rd International Conference on Distributed Computing Systems (ICDCS)*, pp. 83-92, July 2013.
 13. L. Qu, H. Chen, and **Y. Tu**. Nonparametric Copula Estimation in Sensor Networks. In Proceedings of *7th International Conference on Mobile Ad-Hoc and Sensor Networks (MSN)*, pp. 1-8, Dec. 2012.
 14. D. Yu, **Y. Tu**, and H. Yang. Parallel computing simulation of electrical excitation and conduction in a 3D human heart. To appear in *7th INFORMS Workshop on Data Mining and Health Informatics*. Phoenix, AZ, USA. October 2012.
 15. **Y. Tu**, S. Chen, S. Pandit, A. Kumar and V. Grupcev. Efficient SDH Computation In Molecular Simulations Data (short paper). To appear in *ACM Conference on Bioinformatics, Computational Biology, and Biomedicine (ACM-BCB)*. Orlando, FL, USA. October 2012.
 16. J. Ge, Y. Xia, and **Y. Tu**. A Discretization Algorithm for Uncertain Data. In Proceedings of *21st International Conference on Database and Expert Systems Applications (DEXA)*, pp. 485-499.
 17. J. Huang, F. Nie, H. Huang, and **Y. Tu**. Trust Prediction via Aggregating Heterogeneous Social Networks. In Proc. *ACM Conf. Information and Knowledge Management (CIKM)*. pp. 1774-1778, Maui, Hawaii, USA, October 2012.
 18. A. Kumar, V. Grupcev, Y. Yuan, **Y. Tu** and G. Shen. Distance Histogram Computation Based on Spatiotemporal Uniformity in Scientific Databases. In Procs. *International Conference on Extending Database Technology (EDBT)*, pp. 288-299, Berlin, Germany, March 26-30, 2012.
 19. Lixi Wang, Jing Xu, Ming Zhao, Yi-Cheng Tu, and Jose Fortes. Fuzzy Modeling Based Resource Management for Virtualized Database Systems. In Procs. *19th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Communication Systems (MASCOTS)*, pp. 32-42, Singapore, July 25-27, 2011.
 20. **Y. Tu**, X. Wang, and Z. Xu. Power-Aware DBMSs: Potential and Challenges. In Procs. *Scientific and Statistical Database Management Conference (SSDBM)*. pp. 598-599, Portland, WA, July 20-22, 2011.
 21. Z. Xu, **Y. Tu**, and X. Wang. Exploring Power-Performance Tradeoffs in Database Systems. In *Procs. of 26th International Conference on Data Engineering (ICDE)*, pp. 485-496, Long Beach, CA, March 2010.
 22. B. Qin, Y. Xia, R. Sathyesh, S. Prabhakar, and **Y. Tu**. uRule: A Rule-Based Classification System for Uncertain Data (software demo). In Procs. *International Conference on Data Mining (ICDM)*. Miami, FL, December 2009.
 23. **Y. Tu**, S. Chen and S. Pandit. Computing Spatial Distance Histograms Efficiently in Scientific Databases. In Proceedings of *International Conference of Data Engineering (ICDE)*, pp.796-807, Shanghai, China, March 29-April 4, 2009.
 24. B. Qin, Y. Xia, S. Prabhakar, **Y. Tu**. A Rule-based Classification Algorithm for Uncertain Data. In *Proceedings of Workshop of Management and Mining of Uncertain Data*, in conjunction with ICDE'09. pp.1633-1640.
 25. **Y. Tu**, S. Liu, S. Prabhakar, B. Yao, and W. Schroeder. Using Control Theory for Load Shedding in Data Stream Management. In Proceedings of *International Conference on Data Engineering (ICDE)*, pp.1491-1492, Istanbul, Turkey, April 2007.

Yicheng Tu

26. Y. Xia, **Y. Tu**, M. Atallah, and S. Prabhakar. Efficient Data Compression in Location Based Services. In Proceedings of *2nd International Conference on Geosensor Networks*, Boston, MA, October 2006.
27. **Y. Tu**, S. Liu, S. Prabhakar, and B. Yao. Load Shedding in Stream Databases - A Control-Based Approach. In Proceedings of *International Conference on Very Large Databases (VLDB)*, pp.787-798, Seoul, Korea, September 2006.
28. **Y. Tu** and S. Prabhakar. Control-Based Load Shedding in Data Stream Management Systems. PhD Symposium/Workshop, in conjunction with ICDE 2006.
29. **Y. Tu**, J. Yan, and S. Prabhakar. Quality-Aware Replication of Multimedia Data. In Proceedings of *International Conference of Database and Expert Systems Applications (DEXA)* 2005, pp.240-249.
30. **Y. Tu**, M. Hefeeda, Y. Xia, S. Prabhakar, and S. Liu. Control-based Quality Adaptation in Data Stream Management Systems. In Proceedings of *International Conference of Database and Expert Systems Applications (DEXA)* 2005, pp.746-755.
31. R. Cheng, B. Kao, S. Prabhakar, A. Kwan, and **Y. Tu**. Adaptive Stream Filters for Entity-Based Queries with Non-Value Tolerance. Proceedings of *International Conference on Very Large Databases (VLDB)*, pp.37-48.
32. L. Qu and **Y. Tu**. Change Point Estimation of Bar Code Signals. In Proceedings of *International Conference on Scientific Computing (CSC)* 2005, pp.109-114
33. **Y. Tu**, S. Prabhakar, A. Elmagarmid and R. Sion. QuaSAQ: An Approach to Enabling End-to-End QoS for Multimedia Databases. In Proceedings of *International Conference on Extending Database Technology (EDBT)* 2004, pp. 694-711.
34. **Y. Tu**, J. Sun and S. Prabhakar. Performance Analysis of A Hybrid Media Streaming System. In Proceedings of *ACM/SPIE Conference on Multimedia Computing and Networking (MMCN)* 2004, pp. 69-82.
35. W. Aref, A. Catlin, A. Elmagarmid, J. Fan, M. Hammad, I. Ilyas, M. Marzouk, S. Prabhakar, **Y. Tu** and X. Zhu. VDBMS: A Testbed Facility for Research in Video Database Benchmarking. In Proceedings of *International Conference on Distributed Multimedia Systems (DMS)* 2003, pp. 160-166.
36. W. Aref, A. Elmagarmid, J. Fan, J. Guo, M. Hammad, I. Ilyas, M. Marzouk, S. Prabhakar, A. Rezgui, A. Teoh, E. Terzi, **Y. Tu**, A. Vakali, X. Zhu. A Distributed Database Server for Continuous Media. In Proceedings of *International Conference on Data Engineering (ICDE)* 2002, pp. 490-491.

AWARDS AND HONORS

- USF Outstanding Research Achievement Award, 2014
- USF Outstanding Faculty Award, 2014
- NSF CAREER Award, 2013
- New Researcher Award, USF, 2008
- Key member of the team that wins the 2004 MIRA award, a prestigious education/research award in Indiana, for the technical contribution in the Knowledge Projection project.

PROFESSIONAL MEMBERSHIP

- Full member of *Sigma-Xi*, the scientific research society
- Member of *Upsilon Pi Epsilon*, the computer science honorary society
- Member of *ACM*
- Senior Member of *IEEE*
- Member of *ASEE*
- Member of Biophysical Society (BPS)

SERVICES

Journal Editorship

- *Journal of Computational Information Systems*
- *Journal on Knowledge and Data Engineering*

Member of Technical Program Committee

- *IEEE Conference on Networking, Architecture, and Storage (NAS)* 2012
- *IEEE International Conference on Tools with Artificial Intelligence (ICTAI)*, 2011
- *ACM International Conference on Management of Data (SIGMOD)*, demo track, 2010

Yicheng Tu

- *International Conference on Machine Learning and Applications (ICMLA)*, 2010
- *International Conference on Semantics, Knowledge and Grid*, 2010, 2011, 2012
- *Transaction Processing Performance Council Technical Conference (TPC-TC)*, 2009, 2010
- *International Conference on Web Information Systems and Mining (WISM)*, 2008, 2009
- *Software Engineering Research, Management and Applications*, 2009
- *Joint conference of Asia-Pacific Web Conference and Web-Age Information Management*, 2009
- *International Conference on Computational Science and Engineering*, 2008

Journal Reviewer

- *ACM Transactions on Database Systems*
- *IEEE Transactions on Knowledge and Data Engineering*
- *ACM Transactions on Multimedia Computing, Communications, and Applications*
- *IEEE Multimedia*
- *IEEE Transactions on Circuits and Systems for Video Technology*
- *ACM/Springer Multimedia Systems Journal*
- *IEEE Transactions on Systems, Man, and Cybernetics, Part A*
- *Distributed and Parallel Databases*
- *Multimedia Tools and Applications*
- *Data and Knowledge Engineering*
- *Knowledge and Information Systems*
- *Neurocomputing*
- *Journal of Intelligent and Fuzzy Systems*
- *Journal of Information Processing*

Funding Program Reviewer/Panelist

- Panelist in 3 NSF BigData panels, one IIS core program panel, one CRI panel
- Co-Chair, Topics in Computational Biosciences, Center for Scientific Review, National Institutes of Health (NIH), November 2015; panelist on November 2016
- Reviewer, *University of Southern California METRANS* research program, 2009, 2010