

# Jeong-Hyon Hwang

Department of Computer Science  
University at Albany - State University of New York  
Albany, NY 12222  
jhh@cs.albany.edu  
<http://www.cs.albany.edu/~jhh>

Office: (518) 437-3662  
Cell: (401) 431-5229  
Fax: (518) 442-5638

## RESEARCH INTERESTS

---

Databases and Distributed Systems (detailed topics: graph database systems, trajectory data management, real-time data stream processing, Internet-scale data management, fault tolerance and load management)

## EDUCATION

---

<b>Brown University</b>	<b>Providence, RI</b>
Ph.D. in Computer Science	07/2008
Dissertation: Fast and Highly-Available Stream Processing	
Advisor: Professor Stan Zdonik	
M.S. in Computer Science	05/2003
<b>Korea University</b>	<b>Seoul, Korea</b>
M.S. in Computer Science and Engineering	02/2000
B.S. in Computer Science and Engineering	02/1998
B.S. in Mathematics Education	02/1994

## EMPLOYMENT

---

Associate Professor, Department of Computer Science, University at Albany - SUNY	09/2014 - present
Assistant Professor, Department of Computer Science, University at Albany - SUNY	09/2008 - 08/2014
Research Assistant, Department of Computer Science, Brown University	09/2001 - 07/2008
Intern, IBM Thomas J. Watson Research Center, Hawthorne, NY	05/2006 - 09/2006
Software Development Team Manager, On-Korea.com, Seoul, Korea	04/2000 - 08/2001
SW Developer/Soldier, Personnel Management Office, 1st Republic of Korea Army	09/1994 - 11/1996

## HONORS AND AWARDS

---

- Best Poster Award, IEEE International Conference on Data Engineering (ICDE), 04/03/2014
- NSF CAREER award, 01/2012
- Best Poster Presentation Runners Up, ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS), 11/2010
- IBM Open Collaborative Faculty Award, 06/2010
- Best Demo Award, ACM SIGMOD International Conference on Management of Data, 06/2005
- National Scholarship, Ministry of Information and Communication of Korea, 09/2001 - 08/2005
- New Software Award (Grand Prize, Multimedia Contents Area), Ministry of Information and Communication of Korea, 02/2001
- Award for Excellence in Undergraduate Studies, Korea University, 1991, 1992, 1993, 1997
- 3rd Rank (out of approximately 30,000 high school students) in '89 Kangwon Province Mathematics Contest, Korea, 09/1989

## RESEARCH FUNDING

---

### Current

- Sole PI, NSF CAREER award IIS-1149372, ``CAREER: G\*: A Parallel System for Efficiently Processing Large Graphs'', Duration: 02/2012 - 01/2017, Total award: \$496,648

### Completed

- Co-PI (PI: Catherine Lawson, Other Co-PI: S. S. Ravi, Feng Chen), US Department of Transportation, ``Techniques for Information Extraction from Compressed GPS Traces'', Duration: 3/01/2014 - 2/28/2015, Total award: \$100,000

- Sole PI, University at Albany Faculty Research Awards Program (FRAP) - Category A, ``G\*: A Parallel System for Efficiently Managing Large Graphs'', Duration: 05/2012 - 08/2014, Total award: \$9,796
- Co-PI (PI: Catherine Lawson, Other Co-PI: S. S. Ravi), US Department of Transportation, ``Compressing and Querying Multiple GPS Traces for Transportation Planning'', Duration: 12/01/2011 - 12/31/2012, Total award: \$100,000
- Sole PI, KISTI Supercomputing Center, ``Development of Techniques for Managing Large Amounts of Network Data in the Cloud'', Duration: 07/31/2012 - 11/30/2012, Total award: \$42,991
- Sole PI, SnapLogic, ``Development and Evaluation of Parallel Data Processing Techniques'', Duration: 07/15/2011 - 12/15/2011, Total award: \$15,000
- PI (Co-PIs: Jagdish Gangolly, Daniel Rosenkrantz), Albany Foundation - IBM Open Collaboration, ``Financially-Oriented Stream Pattern Detection Language: Specification and Identification'', Duration: 07/15/2010 - 07/15/2011, Total award: \$40,000
- Co-PI (PI: Catherine Lawson, Other Co-PIs: Siwei Lyu, S. S. Ravi), US Department of Transportation, ``Compression and Mining of GPS Trace Data: New Techniques and Applications for Transportation'', Duration: 01/01/2010 - 12/31/2010, Total award: \$99,848
- Co-PI (PI: S. S. Ravi, Other Co-PIs: Jagdish Gangolly, Siwei Lyu, Daniel Rosenkrantz), Albany Foundation - IBM Open Collaboration, ``Patterns of Interest in Financial Data Streams: Specification and Identification'', Duration: 07/15/2009 - 01/15/2010, Total award: \$40,000
- Sole PI, University at Albany Faculty Research Awards Program (FRAP) - Category A, ``Sensor-Based Real-Time Monitoring and Accident Prevention in the UAlbany Tunnel System'', Duration: 05/2009 - 04/2012, Total award: \$8,379

#### **STUDENT ADVISING** (# denotes a female student and + denotes a minority student)

---

##### **Graduated PhD students**

- Paul Olsen Jr., ``Efficient Execution of Top-k Closeness Centrality Queries'', 01/2012 - 07/2016 (First Employment: The College of Saint Rose, Assistant Professor)
- Alan Labouseur, 05/2012 - 05/2014, Thesis: A Query-Oriented Approach to Graph Series Distribution and Replication (Current Employment: Marist College, Assistant Professor)
- Fan Ping, Thesis: Replica Placement using Network Coordinates, 01/2009 - 07/2011 (*First Employment: Amazon*)

##### **Current PhD students**

- Baibhav Rajbhandari, 01/2014 - present
- Aparna Joshi#, 06/2013 - present
- Rohini Vabbalareddy#, 09/2012 - present
- Jayadevan Vijayan, 09/2011 - present
- Jeremy Birnbaum, 05/2011 - present

##### **PhD committees (Graduated)**

- Ewa Musial# (Computer Science, University at Albany - SUNY), Thesis: Effective Entity Resolution Methodology for Improving Data Quality and Reliability of Service-Oriented Applications, 01/2011 - 05/2014
- Steve Lackey (Informatics, University at Albany - SUNY), Thesis: Integrating Place and Time with Tasks: Supporting the Student Commuter, 05/2009 - 05/2014
- Jonathan Muckell (Informatics, University at Albany - SUNY), Thesis: Compression of GPS Trajectory Data: Benchmarking Framework and New Approach, 05/2009 - 05/2013 (*First Employment: NYS Office of Information Technology Services; University at Albany Distinguished Dissertation Award*)
- Xunyu Pan (Computer Science, University at Albany - SUNY), Thesis: Digital Forensics using Local Signal Statistics, 10/2009 - 11/2011 (*First Employment: Assistant Professor, Frostburg State University, MD, USA*)

##### **PhD committees (Current)**

- Aleksandr Krymer (Computer Science, University at Albany - SUNY), 01/2009 - present

##### **Graduated MS students**

- Kyuseo Park, 06/2013 - 05/2014, Project: TrajMetrix: A Trajectory Compression Benchmarking Framework

- Yuchao Ma<sup>#</sup>, 06/2013 - 05/2014, Project: Implementation of Various Single Trajectory Compression Algorithms
- Hsiang-Cheng Meng, Project: Similarity-Based Compression of GPS Trajectory Data, 05/2013 (*First Employment: Tekelec*)
- Vikram Patil, Project: Traffic Congestion Modeling using New York City Taxi Data, 12/2012 (*First Employment: New York State Senate*)
- Kwangjin Gee, Project: JVM Performance Monitoring, 12/2010 (*First Employment: SK C&C*)
- Christopher McConnell, Project: Detouring and Replication for Fast and Reliable Internet-Scale Stream Processing, 05/2010 (*First Employment: GE Research, Graduated with Publication Award, Current Employment: Cloudera*)
- Jerry Lin, Project: Scheduling of Stream Processing Operators over Multiprocessors, 05/2010 (*First Employment: GE Research, Current Employment: Amazon*)
- Gagan Singh, Project: Using the Hadoop Framework for Large Scale Epidemic Simulations, 05/2010 (*First Employment: MSE Power Systems*)
- Paarul Singh<sup>#</sup>, Project: Energy-Efficient Localization in Wireless Sensor Networks, 05/2010
- David Moore Jr., Project: Efficient Processing of GPS Spatial Temporal Data Streams, 05/2010

#### **Current MS students**

- Aparna Joshi<sup>#</sup>, 06/2013 - present
- Daniel Kemp, 02/2012 - present

#### **Past Undergraduate Students**

- Vincent Cerchia (University at Albany), 01/2014 - present
- Christian Ramson<sup>#+</sup> (Dillard University, New Orleans, Louisiana), University at Albany Summer Research Program, 06/2012 - 07/2012
- Daniel Bokser (University at Albany), 09/2011 - 09/2012
- Brandon Jennings<sup>+</sup> (University of Maryland -- Baltimore County), University at Albany Summer Research Program, 06/2010 - 07/2010

#### **Past High School Students**

- Guodong Fu (Montgomery Blair High School, MD), 06/2009 - 08/2009 (*admitted to the University of Maryland -- College Park*)

### **COLLABORATION WITH VISITING SCHOLARS**

---

- Ui-Sung Song, Associate Professor, Busan National University of Education, 08/2015 - 07/2016.
- So-Young Park, Associate Professor, Sangmyung University, 08/2015 - 07/2016.
- Hye-Young Kim, Associate Professor, Hongik University, Korea, 01/2014 - 01/2015.
- Youn-Hee Han, Associate Professor, Korea University of Technology and Education, Korea, 08/2013 - 01/2015.
- Dong-Yong Cho, Professor, Department Computer Science and Engineering, Jeonju University, Korea, 08/2010 - 07/2011.
- Chanyeol Park, Senior Researcher, KISTI (Korea Institute of Science and Technology Information) Super Computing Center, 04/2010 - 03/2011.

**PUBLICATIONS** (overall Google citation count as of 10/02/2013: 1532, \* denotes a student coauthor advised by Jeong-Hyon Hwang)

---

#### **Refereed Journal Publications**

1. Alan Labouseur\*, Paul Olsen Jr.\*, **Jeong-Hyon Hwang**, "Scalable Distribution and Replication of Dynamic Graph Data", under preparation for Proceedings of the VLDB Endowment (PVLDB).
2. **Jeong-Hyon Hwang**, Fan Ping\*, Alan Labouseur\*, "Replica Placement for Fast and Reliable Data Access in Wide Area Networks", under preparation for IEEE Transactions on Parallel and Distributed Systems.
3. Chan-Myung Kim, Yong-Hwan Kim, Youn-Hee Han, and **Jeong-Hyon Hwang**, "Betweenness Centrality Estimation for Social-Aware Routing in Delay-Tolerant Networks", Springer Mobile Networks and Applications (MONET), 21(3): 469-481, 2016.

4. Alan Labouseur\*, Jeremy Birnbaum\*, Paul Olsen Jr.\*, Sean Spillane\*, Jayadevan Vijayan\*, Wook-Shin Han, **Jeong-Hyon Hwang**, "The G\* Graph Database: Efficiently Managing Large Distributed Dynamic Graphs", *Distributed and Parallel Databases (DAPD)*, 33(4): 479-514, 2015.
5. Jonathan Muckell, Paul Olsen Jr.\*, **Jeong-Hyon Hwang**, Catherine Lawson, S. S. Ravi, "Compression of Trajectory Data: A Comprehensive Evaluation and New Approach", *GeoInformatica*, 18(3): 435-460, 2014.
6. SungHo Cho, **Jeong-Hyon Hwang**, Kyung Yul Bae and Chong-Sun Hwang, "An Optimistic Cache Consistency Protocol using a Preemptive Approach", *IEICE Transactions on Information and Systems*, E83-D(9): 1772-1780, 2000.
7. SungHo Cho, **Jeong-Hyon Hwang**, Kyung Yul Bae and Chong-Sun Hwang, "Run and Hit: Optimistic Concurrency Control for Mobile Computing", *Parallel and Distributed Systems and Networks*, 3(3): 171-179, 2000.

#### ***Refereed Conference/Workshop Publications***

8. Alan Labouseur, Justin Svegliato, and **Jeong-Hyon Hwang**, "Distributed Graph Snapshot Placement and Query Performance in a Data Center Environment", to appear in *Proc. of the 2015 International Conference on Computational Science and Computational Intelligence (CSCI)*, Las Vegas, Nevada, 2015.
9. Paul Olsen Jr.\*, Alan Labouseur\*, and **Jeong-Hyon Hwang**, "Efficient Top-k Closeness Centrality Search", *Proc. of the 30th International Conference on Data Engineering (ICDE)*, 196-207, Chicago, Illinois, 2014 (acceptance rate: 19.96% of 446).
10. Alan Labouseur\*, Paul Olsen Jr.\*, Kyuseo Park\*, and **Jeong-Hyon Hwang**, "A Demonstration of Query-Oriented Distribution and Replication Techniques for Dynamic Graph Data", *Proc. of the 23rd World Wide Web Conference (WWW)*, 127-130, Seoul, Korea, 2014 (acceptance rate: 33.3% of 84).
11. Ewa Musial, **Jeong-Hyon Hwang**, Mei-Hwa Chen, and S. S. Ravi, "Efficient Entity Resolution for Heterogeneous Datasets", *Proc. of the 23rd International Conference on Software Engineering and Data Engineering (SEDE)*, 111-118, New Orleans, Louisiana, 2014.
12. Kyuseo Park\*, Jeremy Birnbaum\*, Yuchao Ma\*, Paul Olsen Jr.\*, Jonathan Muckell, S. S. Ravi, Catherine Lawson, and **Jeong-Hyon Hwang**, "TrajMetrix: A Trajectory Compression Benchmarking Framework (Demo Paper)", *Proc. of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS)*, 562-565, Orlando, Florida, 2013.
13. Alan Labouseur\*, Paul Olsen Jr.\*, and **Jeong-Hyon Hwang**, "Scalable and Robust Management of Dynamic Graph Data", *Proc. of the First International Workshop on Big Dynamic Distributed Data (BD<sup>3</sup>)*, held in conjunction with the 39th International Conference on Very Large Data Bases (VLDB), 43-48, Riva del Garda, Italy, 2013.
14. Mert Akdere, **Jeong-Hyon Hwang**, and Ugur Cetintemel, "Real-time Probabilistic Data Association over Streams", *Proc. of the 7th ACM International Conference on Distributed Event-Based Systems (DEBS)*, 219-230, Arlington, Texas, 2013 (acceptance rate: 27.5% of 58).
15. Jeremy Birnbaum\*, Hsiang-Cheng Meng\*, **Jeong-Hyon Hwang**, and Catherine Lawson, "Similarity-Based Compression of GPS Trajectory Data", *Proc. of the 4th International Conference on Computing for Geospatial Research and Application (COM.Geo)*, 92-95, San Jose, California, 2013.
16. Jonathan Muckell, Paul Olsen Jr.\*, **Jeong-Hyon Hwang**, S. S. Ravi, Catherine Lawson, "A Framework for Efficient and Convenient Evaluation of Trajectory Compression Algorithms", *Proc. of the 4th International Conference on Computing for Geospatial Research and Application (COM.Geo)*, 24-31, San Jose, California, 2013.
17. Sean Spillane\*, Jeremy Birnbaum\*, Daniel Bokser\*, Daniel Kemp\*, Alan Labouseur\*, Paul Olsen Jr.\*, Jayadevan Raja\*, **Jeong-Hyon Hwang**, and Jun-Weon Yoon, "A Demonstration of the G\* Graph Database Systems", *Proc. of the 29th International Conference on Data Engineering (ICDE)*, 1356-1359, Brisbane, Australia, 2013.
18. **Jeong-Hyon Hwang**, Jeremy Birnbaum\*, Rohini Vabbalareddy\*, S. S. Ravi, and Chanyeol Park, "A Graph Database Approach for Efficient and Scalable Management of Simulations", *Proc. of the 2nd Annual Workshop on High-Performance Computing meets Databases (HPCDB)*, held in conjunction with the Conference on High Performance Computing Networking, Storage and Analysis (SC), 1310-1311, Salt Lake City, Utah, 2012.
19. Jonathan Muckell, Vikram Patil\*, Fan Ping\*, **Jeong-Hyon Hwang**, Catherine Lawson, and S. S. Ravi, "SQUISH: An Online Approach for GPS Trajectory Compression", *Proc. of the 2nd International*

- Conference on Computing for Geospatial Research and Application (COM.Geo), 13.1-13.8, Washington, DC, 2011 (acceptance rate: 33.8% of 157).
20. Fan Ping\*, Xiaohu Li\*, Christopher McConnell\*, Rohini Vabbalareddy\*, and **Jeong-Hyon Hwang**, "Towards Optimal Data Replication Across Data Centers", Proc. of the 1st Workshop on Data Center Performance (DCPerf), held in conjunction with the 31st International Conference on Distributed Computing Systems (ICDCS), 66-71, Minneapolis, Minnesota, 2011.
  21. Jonathan Muckell, **Jeong-Hyon Hwang**, Catherine Lawson, and S. S. Ravi, "Algorithms for Compressing GPS Trajectory Data: An Empirical Evaluation", Proc. of 18th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS), 402-405, San Jose, California, 2010 (**best poster presentation runners up**).
  22. Christopher McConnell\*, Fan Ping\*, and **Jeong-Hyon Hwang**, "iFlow: An Approach for Fast and Reliable Internet-Scale Stream Processing Utilizing Detouring and Replication", Proc. of the 36th International Conference on Very Large Data Bases (VLDB), 1557-1560, Singapore, 2010.
  23. Fan Ping\*, Christopher McConnell\*, and **Jeong-Hyon Hwang**, "A Retrospective Approach for Accurate Network Latency Prediction", Proc. of the 2nd Workshop on Grid and P2P Systems and Applications (GridPeer), held in conjunction with the 19th International Conference on Computer Communications and Networks (ICCCN), 1-6, Zurich, Switzerland, 2010.
  24. Christopher McConnell\*, Fan Ping\*, and **Jeong-Hyon Hwang**, "Detouring and Replication for Fast and Reliable Internet-Scale Stream Processing", Proc. of the 3rd International Workshop on Data Intensive Distributed Computing (DIDC), held in conjunction with the 19th International Symposium on High Performance Distributed Computing (HPDC), 737-745, Chicago, Illinois, 2010.
  25. **Jeong-Hyon Hwang**, Sanghoon Cha, Ugur Cetintemel, and Stan Zdonik, "Borealis-R: A Replication-Transparent Stream Processing System for Wide-Area Monitoring Applications (Demonstration)", Proc. of the ACM SIGMOD International Conference on Management of Data, 1303-1306, Vancouver, Canada, 2008.
  26. **Jeong-Hyon Hwang**, Ugur Cetintemel, and Stan Zdonik, "Fast and Highly-Available Stream Processing over Wide Area Networks", Proc. of the 24th International Conference on Data Engineering (ICDE), 804-813, Cancun, Mexico, 2008 (acceptance rate: 12.1% of 617).
  27. **Jeong-Hyon Hwang**, Ying Xing, Ugur Cetintemel, and Stan Zdonik, "A Cooperative, Self-Configuring High-Availability Solution for Stream Processing", Proc. of the 23rd International Conference on Data Engineering (ICDE), 176-185, Istanbul, Turkey, 2007 (acceptance rate: 18.5% of 659).
  28. **Jeong-Hyon Hwang**, Ugur Cetintemel, and Stan Zdonik, "Fast and Reliable Stream Processing over Wide Area Networks", Proc. of the 1st IEEE Workshop on Scalable Stream Processing systems (SSPS), held in conjunction with the 23rd International Conference on Data Engineering (ICDE), Istanbul, Turkey, 2007.
  29. Ying Xing, **Jeong-Hyon Hwang**, Ugur Cetintemel, and Stan Zdonik, "Providing Resiliency to Load Variations in Distributed Stream Processing", Proc. of the 32nd International Conference on Very Large Data Bases (VLDB), 775-786, Seoul, Korea, 2006 (acceptance rate: 13.9% of 331).
  30. Y. Ahmad, U. Cetintemel, B. Berg, M. Humphrey, **J.-H. Hwang**, O. Papaemmanouil, A. Rasin, N. Tatbul, Y. Xing, and S. Zdonik, "Distributed Operation in the Borealis Stream Processing Engine (Invited Demonstration)", Proc. of the 2nd International Conference on Geosensor Networks, Boston, Massachusetts, 2006.
  31. Y. Ahmad, U. Cetintemel, B. Berg, M. Humphrey, **J.-H. Hwang**, O. Papaemmanouil, A. Rasin, N. Tatbul, Y. Xing, and S. Zdonik, "Distributed Operation in the Borealis Stream Processing Engine (Demonstration)", Proc. of the ACM SIGMOD International Conference on Management of Data, 882-884, Baltimore, Maryland, 2005 (**best demo award**).
  32. **Jeong-Hyon Hwang**, Magdalena Balazinska, Alexander Rasin, Ugur Cetintemel, Michael Stonebraker, and Stan Zdonik, "High-Availability Algorithms for Distributed Stream Processing", Proc. of the 21st International Conference on Data Engineering (ICDE), 779-790, Tokyo, Japan, 2005 (acceptance rate: 12.9% of 521).
  33. Ying Xing, Stan Zdonik, and **Jeong-Hyon Hwang**, "Dynamic Load Distribution in the Borealis Stream Processor", Proc. of the 21st International Conference on Data Engineering (ICDE), 791-802, Tokyo, Japan, 2005 (acceptance rate: 12.9% of 521).
  34. M. Balazinska, D. Abadi, Y. Ahmad, M. Cherniack, **J.-H. Hwang**, W. Lindner, A. Rasin, N. Tatbul, Y. Xing, S. Zdonik, "The Design of the Borealis Stream Processing Engine", Proc. of 2nd Biennial Conference on Innovative Data Systems Research (CIDR), 277-289, Asilomar, California, 2005.
  35. D. Abadi, D. Carney, U. Cetintemel, M. Cherniack, C. Convey, C. Erwin, E. Galvez, M. Hatoun, **J.-H. Hwang**, A. Maskey, A. Rasin, A. Singer, M. Stonebraker, N. Tatbul, Y. Xing, R. Yan, S. Zdonik, "Aurora:

- A Data Stream Management System (Demonstration)", Proc. of the ACM SIGMOD International Conference on Management of Data, 666, San Diego, California, 2003.
36. **Jeong-Hyon Hwang**, SungHo Cho, and Chong-Sun Hwang, "Optimized Scheduling on Broadcast Disks", Proc. of the 2nd International Conference on Mobile Data Management (MDM), 91-104, Hong Kong, China, 2001.
  37. **Jeong-Hyon Hwang**, SungHo Cho, and Chong-Sun Hwang, "Selection of Disk Frequencies for Optimal Multidisk Broadcast Program Generation", Proc. of the 26th Korea Information Science Society Fall Conference, 26 (3): 252-254, Seoul, Korea, 1999.
  38. SungHo Cho, Kyung Y. Bae, **Jeong-Hyon Hwang**, and Chong-Sun Hwang, "An Optimistic Two-Phase Locking Method with a No-Wait Approach", Proc. of HPC Asia 1998, 1414-1421, Singapore, 1998.
  39. **Jeong-Hyon Hwang**, Joung-Hoon Kim, Jin-Ho Ahn, and Chong-Sun Hwang, "Implementation of a Ray-Tracer on a Java-based Parallel-Execution Environment", Proc. of the 25th Korea Information Science Society Fall Conference, 25 (2): 771-773, Seoul, Korea, 1998.

#### **Book Chapters**

40. **Jeong-Hyon Hwang**, Alan Labouseur\*, Paul Olsen\*, "Text Stream Processing", to appear in Encyclopedia of Database Systems, Ling Liu, Tamer Ozsu (editors), Springer.
41. **Jeong-Hyon Hwang**, Alan Labouseur\*, Paul Olsen\*, "Non-Relational Streams", to appear in Encyclopedia of Database Systems, Ling Liu, Tamer Ozsu (editors), Springer.
42. Magdalena Balazinska, **Jeong-Hyon Hwang**, Mehul Shah, "High Availability and Fault Tolerance in Data Stream Management Systems", Encyclopedia of Database Systems, Ling Liu, Tamer Ozsu (editors), Springer (ISBN: 978-0-387-49616-0), 1109-1115, 2009.
43. N. Tabul, Y. Ahmad, U. Cetintemel, **J.-H. Hwang**, Y. Xing, S. Zdonik, "Load Management and High Availability in the Borealis Distributed Stream Processing Engine", Advances in Geosensor Networks, S. Nittel, A. Labrinidis, A. Stefanidis (editors), Springer-Verlag (ISBN: 978-3-540-79995-5), 66-85, 2008.
44. U. Cetintemel, D. Abadi, Y. Ahmad, H. Balakrishnan, M. Balazinska, M. Cherniack, **J.-H. Hwang**, W. Lindner, S. Madden, A. Maskey, A. Rasin, E. Ryvkina, M. Stonebraker, N. Tatbul, Y. Xing, S. Zdonik, "The Aurora and Borealis Stream Processing Engines", Data Stream Management: Processing High-Speed Data Streams, M. Garofalakis, J. Gehrke, R. Rastogi (editors), Springer-Verlag, 2006.

#### **Patents**

- Chitra Dorai, **Jeong-Hyon Hwang**, Rob Strom, "Stochastic Control Optimization for Sender-based Flow Control in a Distributed Stateful Messaging System", US Patent 7761401, 2010.

#### **Books**

- **Jeong-Hyon Hwang**, "Algorithms in C (Korean translation of a book originally written by Robert Sedgewick)", 800 pages, Pearson Korea, 2002.
- **Jeong-Hyon Hwang**, "Java2 Unleashed (Korean translation of a book originally written by Jamie Jaworski)", 1200 pages, Daerim pub. Co., 1999.

#### **Open-Source Software**

- G\*: a parallel system for efficiently storing and querying series of graphs that represent evolving networks, supported by NSF CAREER award IIS-1149372, approximately 20,000 lines of source code, <http://www.cs.albany.edu/~gstar/>

#### **PRESENTATIONS**

- "Efficient Management of Historic Graph Data and High-Volume Data Streams", Pohang University of Science and Technology (POSTECH), Pohang, Korea, 07/17/2014.
- "Efficient Management of Historic Graph Data and High-Volume Data Streams", Rational Retention, Albany, NY, 04/03/2014.
- "Efficient Management of Historic Graph Data and High-Volume Data Streams", Kitware, Clifton Park, NY, 12/17/2013.
- "Scalable and Robust Management of Dynamic Graph Data", BD<sup>3</sup> 2013, Riva del Garda, Italy, 08/30/2013.
- "Real-time Probabilistic Data Association over Streams", DEBS 2013, Arlington, Texas, 07/03/2013.
- "Quickly Finding the  $k$  Most Central Entities in Large Networks", NEDB 2013, MIT, Cambridge, Massachusetts, 02/01/2013.

- “G\*: A Parallel System for Efficiently Managing Large Graphs”, NEDB 2012, MIT, Cambridge, Massachusetts, 02/03/2012.
- “Towards Optimal Data Replication Across Data Centers”, DCPeF 2011, Minneapolis, Minnesota, 06/24/2011.
- “Internet-Scale Data Management”, GE Global Research Center, Niskayuna, New York, 11/30/2010.
- “A Retrospective Approach for Accurate Network Latency Prediction”, GridPeer 2010, ETH Zurich, Zurich, Switzerland, 08/05/2010.
- “Fast and Highly-Available Stream Processing over Wide Area Networks”, US-Korea Conference on Science, Technology, and Entrepreneurship, Raleigh, North Carolina, 07/17/2009.
- “iFlow: An Internet-Scale Data Processing System”, Virginia Bioinformatics Institute, Virginia Tech, Blacksburg, Virginia, 07/16/2009.
- “iFlow: An Internet-Scale Software System for Data-Intensive Computing”, New Trends in Informatics Research, State University of New York at Albany, New York, 04/17/2009.
- “Fast and Highly-Available Stream Processing over Wide Area Networks”, ICDE 2008, Cancun, Mexico, 04/10/2008.
- “Fast and Highly-Available Stream Processing”, IBM Research, Almaden, California, 05/12/2008.
- “-----”, Department of Computer Science, Texas State University, San Marcos, Texas, 05/09/2008.
- “-----”, Department of Computer Science, State University of New York, Albany, New York, 04/29/2008.
- “-----”, Fuji + Xerox Palo Alto Labs, Palo Alto, California, 04/24/2008.
- “-----”, Department of Mathematical Sciences and Applied Computing, Arizona State University, Phoenix, Arizona, 04/18/2008.
- “-----”, Department of Computer Science, Michigan State University, Lansing, Michigan, 04/04/2008.
- “-----”, VMKernel Group, VMWare, Palo Alto, California, 03/17/2008.
- “-----”, Department of Computer Science, University of Missouri, Rolla, Missouri, 03/04/2008.
- “-----”, AT&T Labs, Florham Park, New Jersey, 02/19/2008.
- “-----”, Department of Computer and Information Sciences, Temple University, Philadelphia, Pennsylvania, 02/13/2008.
- “-----”, NEC Labs, Princeton, New Jersey, 12/12/2007.
- “-----”, Exploratory Stream Processing Systems Group, IBM Thomas J. Watson Research Center, Hawthorne, NY, USA, 05/16/2007.
- “A Cooperative, Self-Configuring High-Availability Solution for Stream Processing”, ICDE 2007, Istanbul, Turkey, 04/17/2007.
- “Fast and Reliable Stream Processing over Wide Area Networks”, SSPS 2007, Istanbul, Turkey, 04/16/2007.
- “Issues in Distributed Stream Processing: Load Management, High Availability, and Others”, KAIST, Daejeon, Korea, 09/22/2006.
- “Introduction to Distributed Stream Processing”, Ubiquitous Computing Lab, IBM Korea, Seoul, Korea, 09/19/2006.
- “Providing Resiliency to Load Variations in Distributed Stream Processing”, VLDB 2006, Seoul, Korea, 09/15/2006.
- “Send or Withhold? Finding an Optimal Flow Control Policy in Stream Processing Systems”, Distributed Messaging Group, IBM T. J. Watson Research Center, Hawthorne, New York, USA, 08/25/2006.
- “Things to Consider When Comparing Stream Processing Engines”, Distributed Messaging Group, IBM Thomas J. Watson Research Center, Hawthorne, New York, USA, 06/06/2006.
- “High-Availability Algorithms for Distributed Stream Processing”, ICDE 2005, Tokyo, Japan, 04/08/2005.
- “Issues in Distributed Stream Processing”, Internet Server Group, Digital Home Division, ETRI, Daejeon, Korea, 10/07/2004.
- “Optimized Scheduling on Broadcast Disks”, MDM 2001, Hong Kong, China, 01/08/2001.

## SCHOLARLY ACTIVITIES

---

### **Chair/Organizer**

- Proceedings Co-Chair, the 31st International Conference on Data Engineering (ICDE 2015)
- Organizer, 2014 ACM SIGMOD Programming Contest
- Publication Chair, the 13th Technical Symposium of the Korean Computer Scientists and Engineers Association in America (KOCSEA 2012)
- Session Chair, the 31st International Conference on Data Engineering (ICDE 2015)
- Session Chair, the 30th International Conference on Data Engineering (ICDE 2014)
- Session Chair, the 31st International Conference on Distributed Computing Systems (ICDCS 2011)

### **Program Committee Member**

- The 2016 ACM SIGMOD International Conference on Management of Data (SIGMOD 2016) Demonstration Track
- The 39th International Conference on Very Large Data Bases (VLDB 2015)
- The 31st International Conference on Data Engineering (ICDE 2015)
- The 2014 ACM SIGMOD International Conference on Management of Data (SIGMOD 2014)
- The 38th International Conference on Very Large Data Bases (VLDB 2014)
- The International Conference on Data Engineering (ICDE) 2014 PhD Symposium
- The 16th Asia-Pacific Web Conferences (APWeb 2014)
- The 4th International Workshop on Data Center Performance (DCPerf 2014), co-located with ICDCS 2014
- The SDM 2014 Workshop on Mining Networks and Graphs (MNG 2014)
- The 37th International Conference on Very Large Data Bases (VLDB 2013)
- The 29th International Conference on Data Engineering (ICDE 2013) Best Poster Award Committee
- The 25th International Conference on Scientific and Statistical Database Management (SSDBM 2013) Best Paper Award Committee
- The 25th International Conference on Scientific and Statistical Database Management (SSDBM 2013)
- The 15th Asia-Pacific Web Conferences (APWeb 2013)
- The 3rd International Workshop on Data Center Performance (DCPerf 2013), co-located with ICDCS 2013
- The 2012 ACM SIGMOD International Conference on Management of Data (SIGMOD 2012) Best Demonstration Award Committee
- The 2012 ACM SIGMOD International Conference on Management of Data (SIGMOD 2012) Demonstration Track
- The 36th International Conference on Very Large Data Bases (VLDB 2012)
- The 1st International Conference on Computer Convergence Technology (ICCT 2011)
- The 3rd International Workshop on Specialized Ad Hoc Networks and Systems (SAHNS 2011), co-located with ICDCS 2011
- The 5th ACM International Conference on Distributed Event-Based Systems (DEBS 2011)
- The 27th International Conference on Data Engineering (ICDE 2011)
- The 2nd International Workshop on Ubiquitous Computing and Applications (IWUCA 2010)
- The 3rd International Workshop on Scalable Stream Processing Systems (SSPS 2010), co-located with IPDPS 2010
- The 12th Asia-Pacific Web Conferences (APWeb 2010)
- The 2nd International Workshop on Specialized Ad Hoc Networks and Systems (SAHNS 2009), co-located with ICDCS 2009
- The 6th International Workshop on Data Management for Sensor Networks (DMSN 2009), co-located with VLDB 2009
- The 11th International Conference on Extending Database Technology (EDBT 2008)

### **Journal Reviewer** (each number in [] denotes the number of papers reviewed in the corresponding year)

- VLDB Journal, 2015, 2013[3], 2010
- ACM Transactions on Database Systems (TODS), 2013



- ACM SIGMOD Record, 2014
- ACM Transactions on Internet Technology (TOIT), 2015
- IEEE Transactions on Knowledge and Data Engineering (TKDE), 2011, 2010, 2009
- IEEE Transactions on Reliability (TR), 2012
- IEEE Transactions on Parallel and Distributed Systems (TPDS), 2011
- Springer Distributed and Parallel Databases (DAPD), 2013, 2010
- Elsevier Data and Knowledge Engineering (DKE), 2013
- Elsevier Information Systems (IS), 2010
- Elsevier Parallel Computing (PARCO), 2009
- Software - Practice and Experience (SPE), 2012, 2010
- International Journal of Parallel, Emergent and Distributed Systems (IJPEDS), 2014
- Journal of Computer Science and Technology (JCST), 2010
- Journal of Information Processing Systems (JIPS), 2010

#### **Government Review Panels**

- NSF Proposal Review Panel, April, 2016
- NSF Proposal Review Panel, June, 2015
- NSF Proposal Review Panel, May, 2013
- NSF Proposal Review Panel, April, 2012
- NSF Proposal Review Panel, March, 2012

#### **External Conference Reviewer**

- The 28th International Conference on Data Engineering (ICDE 2012)
- The 35th International Conference on Very Large Data Bases (VLDB 2009)
- The 2008 ACM SIGMOD International Conference on Management of Data (SIGMOD 2008)
- The 24th International Conference on Data Engineering (ICDE 2008)
- The 38th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN 2008)
- The 33rd International Conference on Very Large Data Bases (VLDB 2007)
- The 1st International Workshop on Distributed Event Processing, Systems and Applications (DEPSA 2007)
- The 26th IEEE International Conference on Distributed Computing Systems (ICDCS 2006).

#### **COLLABORATORS (2006 - present)**

---

Daniel Abadi (Yale University), Yanif Ahmad (Johns Hopkins University), Mert Akdere (Google), Magdalena Balazinska (University of Washington), Petko Bogdanov (University at Albany - SUNY), Feng Chen (University at Albany - SUNY), Ugur Çetintemel (Brown University), Mitch Cherniack (Brandeis University), Wei Fan (IBM Research), Jagdish Gangolly (University at Albany - SUNY), Youn-Hee Han (Korea University of Technology and Education), Wook-Shin Han (Pohang University of Science and Technology, Korea), Catherine Lawson (Department of Geography and Planning, University at Albany - SUNY), Siwei Lyu (University at Albany - SUNY), Olga Papaemmanouil (Brandeis University), Chanyeol Park (KISTI Supercomputing Center, Korea), Sekharipuram S. Ravi (University at Albany - SUNY), Daniel Rosenkrantz (University at Albany - SUNY), Mehul A. Shah (HP Labs), Michael Stonebraker (MIT), Nesime Tatbul (Intel Labs), Octavian Udrea (IBM Research), Jun-Weon Yoon (KISTI Supercomputing Center, Korea), Stan Zdonik (Brown University)

#### **COURSES TAUGHT**

##### **2016**

CSI 508. Database Systems I (Fall)

##### **2015**

CSI 508. Database Systems I (Fall)  
CSI 508. Database Systems I (Spring)

- CSI 445/660. Topics in Data Management Systems (Spring)
- 2014**
- CSI 405. Object Oriented Programming Principles and Practice (Fall)  
 CSI 508. Database Systems I (Spring)  
 CSI 445/660. Topics in Data Management Systems (Spring)
- 2013**
- CSI 405. Object Oriented Programming Principles and Practice (Fall)  
 CSI 508. Database Systems I (Spring)  
 CSI 445/660. Topics in Data Management Systems (Spring)
- 2012**
- CSI 508. Database Systems I (Spring)  
 CSI 445/660. Topics in Data Management Systems (Spring)
- 2011**
- CSI 405. Object Oriented Programming Principles and Practice (Fall)  
 CSI 508. Database Systems I (Spring)  
 CSI 445/600. Distributed Systems (Spring)
- 2010**
- CSI 445/660. Topics in Data Management Systems (Fall)  
 CSI 508. Database Systems I (Spring)  
 CSI 445/600. Distributed Systems (Spring)
- 2009**
- CSI 445/660. Topics in Data Management Systems (Fall)  
 CSI 445/600. Distributed Systems (Spring)
- 2008**
- CSI 445/660. Topics in Data Management Systems (Fall)

## COURSE DEVELOPMENT

---

### ***CSI 445/660. Topics in Data Management Systems***

In 2008, I created a new undergraduate/graduate course that reviews recent trends in the design and implementation of data management systems. Students taking this course present and discuss research articles published in prestigious conferences and journals on database systems and carry out team projects of their own choosing. For these projects, students can use a worldwide network testbed (<http://www.planet-lab.org>), a 72-core server cluster, and a sensor network kit. In each year, usually one or two team projects contributed to academic publications.

### ***CSI 508. Database Systems I***

In 2010, I redesigned a graduate level course on database management systems. This course aims to introduce students to the foundations of database systems with emphasis on relational algebra, query processing, query optimization, and concurrency control. Students taking this course need to complete 6 homework assignments as well as a series of 3-4 programming assignments that typically implement a data storage manager, a B+-tree index, a query processor, and a transaction manager.

### ***CSI 405. Object Oriented Programming Principles and Practice***

In 2011, I redesigned an undergraduate level course that covers object oriented software design principles (abstraction, inheritance, and design patterns) with emphasis on how they are embodied in the Java programming language. Based on my own experience of designing and implementing large scale data management systems, I created lectures on advanced features of Java including reflection, threading, network programming, and remote method invocation. In addition, I developed a new series of 10 programming assignments that eventually develop an online multi-player video game while helping students acquire skills and knowledge for solving various computational problems.

**CSI 445/600. Distributed Systems**

In 2009, I developed an undergraduate/graduate course on distributed systems. I created lectures that gave an overview of distributed systems and provided recent research papers for students to present and discuss in class. I also helped students conduct team projects of their own choosing and present their project outcomes as technical reports. This course was offered in 2009, 2010, and 2011.

**DEPARTMENTAL SERVICE**

---

- Undergraduate Advisor, 09/2009 - present
- Member, Graduate Program Committee, 09/2009 - present
- Member, PhD Analytic Exam Committee, 08/2014
- Member, Computing Oversight Committee, 09/2009 - present
- Member, Faculty Search Committee, 09/2013 - 5/2014
- Member, PhD Analytic Exam Committee, 08/2013
- Member, Faculty Search Committee, 10/2012 - 05/2013
- Co-Author, SUNY 2020 Faculty Hiring Proposal (approved), 08/2012 - 11/2012
- Co-Author, SUNY 2020 Faculty Hiring Proposal (approved), 02/2012 - 05/2012
- Member, Department Chair Selection Committee, 03/2012 - 04/2012
- Member, PhD Analytic Exam Committee, 12/2011
- Member, PhD Analytic Exam Committee, 01/2011
- Member, PhD Analytic Exam Committee, 05/2010
- Member, Department Chair Selection Committee, 12/2009 - 02/2010

**COLLEGE AND UNIVERSITY SERVICE**

---

- Participant, College of Computing and Information Commencement, 05/14/2015
- Participant, College of Computing and Information Commencement, 05/17/2015
- Participant, College of Computing and Information Show Case, 04/29/2015
- Participant, Academic/Student Services Fair, 03/04/2015
- Member, Computer Engineering Faculty Search Committee, 07/2014 - present
- Participant, College of Computing and Information Commencement, 05/17/2014
- Member, Computer Engineering Project Team, 10/2013 - present
- Member, Dialog in Action Series for improving diversity and inclusion at UAlbany, 10/2013 - 04/2014
- Poster Presenter, A Celebration of Research, Scholarship and Creative Activities at UAlbany, 09/25/2013
- Participant, College of Computing and Information Commencement, 05/18/2013
- Volunteer, University at Albany Summer Research Program (UASRP), 06/2012 - 07/2012
- Participant, University at Albany Summer Research Program MD/PhD Reception, 06/28/2012
- Participant, College of Computing and Information Commencement, 05/19/2012
- Volunteer, UAlbany Candlelighting, 08/25/2011
- Participant, College of Computing and Information Commencement, 05/14/2011
- Volunteer, Accepted Student Open House, 04/02/2011
- Participant, Reception for International Visiting Researchers, Scholars, and New Sponsored Students, 10/06/2010.
- Volunteer, UAlbany Candlelighting, 08/26/2010
- Volunteer, University at Albany Summer Research Program (UASRP), 06/2010 - 07/2010
- Participant, College of Computing and Information Commencement, 05/15/2010
- Volunteer, Meet UAlbany Faculty, 10/10/2009
- Speaker, Explore UAlbany Day, 08/28/2009
- Participant, College of Computing and Information Commencement, 05/16/2009
- Volunteer, Accepted Student Open House, 04/18/2009

**MISCELLANEOUS**

---

- Citizenship: South Korea (Permanent Resident of the United States)
- Languages: Korean, English
- Teaching Certificate I, Sheridan Center for Teaching and Learning in Higher Education, Brown University, May 2007