Yanif Ahmad <yna@cs.brown.edu>

Dept. of Computer Science.
4130 Upson Hall, Cornell University.
Ithaca, NY 14853-7501, USA.

Tel: 401-497-2105
Fax: 401-863-7657
Web: http://www.cs.cornell.edu/~yanif

Research Interests

Data management, in particular, data stream processing, numerical and scientific query processing, model-based databases, incremental view maintenance, distributed data management, databases in the cloud and NoSQL-databases, approximate query processing, and constraint databases.

Current Employment

February 2009 - present.

Postdoctoral Associate.
Database Group, Computer Science Department.
Cornell University, Ithaca, NY, USA.
Project: DBToaster: Compiling Main-Memory Database Query Processors.
Mentor: Christoph Koch.

Education

2002 - 2009 Doctor of Philosophy. Computer Science.
Brown University, Providence, RI, USA.
Thesis: Pulse: Database Support for Efficient Query Processing Of Temporal Polynomial Models.
Advisor: Prof. Uğur Çetintemel.
Committee: Stanley B. Zdonik, John Jannotti, Samuel Madden.

2001 - 2004 Master of Science. Computer Science.
Brown University, Providence, RI, USA.
Thesis: SAND: Scalable Adaptive Network Databases.
Advisor: Prof. Uğur Çetintemel.
Committee: Stanley B. Zdonik, Eli Upfal.

1998 - 2001 B.Eng. Computing.
Imperial College of Science, Technology and Medicine, London, UK.
Thesis: CVE - A Collaborative Virtual Environment.
Advisor: Prof. John Darlington.

Awards

IBM Ph.D. Fellowship, 2008-2009.
ICDE Best Research Paper, 2008.
SIGMOD Best Demonstration, 2005.
Brown University Fellowship, 2002.

Research Experience

Intern, Computer Science Group, IBM Almaden Research Center. Summer 2008.
Mentor: Sandeep Tata. Manager: Latha Colby, Guy Lohman.
Investigated the viability of traditional database query processors for performing streaming processing, side-by-side with existing enterprise deployments of DBMS.

Intern, Networked Embedded Computing Group, Microsoft Research. Summer 2006.
Mentor: Suman Nath. Manager: Feng Zhou.
Designed and prototyped a communication-efficient index structure for a sensor web portal, exploiting caching and sampling techniques for reduced sensor probing (c.f. COLR-Tree).

Conference and Workshop Publications

DBToaster: Compiling Main-Memory Database Query Processors.
Yanif Ahmad, Christoph Koch.
Under submission.

Exact Online Aggregation by Message Passing in Main-Memory Cloud Warehouses.
Oliver Kennedy, Christoph Koch, Yanif Ahmad.
Under submission.

Simultaneous Equation Systems for Query Processing on Continuous-Time Data Streams.
Yanif Ahmad, Olga Papaemmanouil, Uğur Çetintemel, Jennie Rogers.
Proceedings of the 24th International Conference on Data Engineering (ICDE '08).

COLR-Tree: Communication Efficient Spatio-Temporal Index for a Sensor Data Web Portal.
Yanif Ahmad, Suman Nath.
Proceedings of the 24th International Conference on Data Engineering (ICDE '08).
(Best Research Paper Award)

Declarative Temporal Data Models for Sensor-Driven Query Processing.
Yanif Ahmad, Uğur Çetintemel.
Proceedings of the International Workshop on Data Management in Sensor Nets (DMSN '07).

XPORT: Extensible Profile-Drive Overlay Routing Trees.
Olga Papaemmanouil, Yanif Ahmad, Uğur Çetintemel, John Jannotti.
Proceedings of the 25th ACM SIGMOD Conference (SIGMOD '06).

Locality-Aware Networked Join Evaluation.
Yanif Ahmad, Uğur Çetintemel, John Jannotti, Alexander Zgolinski.
Proceedings of the IEEE International Workshop Networking Meets Databases (NetDB '05).

The Design of the Borealis Stream Processing Engine.
with the Borealis Team.
Proceedings of the 2nd Biennial Conference on Innovative Database Systems (CIDR '05).

Network-Aware Query Processing for Stream Based Applications.
Yanif Ahmad, Uğur Çetintemel.
Proceedings of the 30th International Conference on Very Large Data Bases (VLDB '04).

A Type System for Statically Detecting Spreadsheet Errors.
Yanif Ahmad, Tudor Antoniu, Sharon Goldwater, Shriram Krishnamurthi.
Proceedings of the IEEE International Conference on Automated Software Engineering (ASE '03).

Exploiting Precision vs. Efficiency Tradeoffs in Symmetric Replication Environments.
Uğur Cetintemel, Peter Keleher, Yanif Ahmad. (abstract)
Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC '02).

Demonstrations

DBToaster: A SQL Compiler for High-Performance Delta Processing in Main-Memory Databases.
Yanif Ahmad, Christoph Koch.
Proceedings of the VLDB Endowment, Volume 2 (VLDB '09).

The Borealis Distributed Stream Processing Engine.
with the Borealis Team.
Proceedings of the 2nd International Conference on GeoSensor Networks (GSN '06, invited).

XPORT: Extensible Profile-driven Overlay Routing Trees.
Olga Papaemmanouil, Yanif Ahmad, Uğur Çetintemel, John Jannotti, Yenel Yildirim.
Proceedings of the 25th ACM SIGMOD Conference (SIGMOD '06).

Distributed Operation in the Borealis Stream Processing Engine.
with the Borealis Team.
Proceedings of the 24th ACM SIGMOD Conference (SIGMOD '05).
(Best Demonstration Award)

Book Chapters and Articles

1. Data streams: Architectures and prototypes.
2. Data streams: Streaming applications.

Yanif Ahmad, Uğur Çetintemel.
Encyclopedia of Database Systems.
L. Liu, M. Tamer Ozsu (editors), Springer.

Load Management and High Availability in the Borealis Stream Processing Engine.
Nesime Tatbul, Yanif Ahmad, Ugur Cetintemel, Jeong-Hyon Hwang, Ying Xing, Stan Zdonik.
Advances in Geosensor Networks.
S. Nittel, A. Labrinidis, A. Stefanidis (editors), LNCS, Springer-Verlag.

The Aurora and Borealis Stream Processing Engines.
with the Borealis Team.
Data Stream Management: Processing High-Speed Data Streams.
M. Garofalakis, J. Gehrke, R. Rastogi (editors), Springer-Verlag, July 2006.

Network Awareness in Internet Scale Stream Processing.
Yanif Ahmad, Uğur Çetintemel, John Jannotti, Alexander Zgolinski, Stan Zdonik.
IEEE Bulletin of the Technical Committee on Data Engineering, March 2005, Vol 28, No. 1.

Patents

An Aggressive Compilation Framework For Continuous Query Processing on Update Streams.
Yanif Ahmad, Christoph Koch.
US Provisional Patent Application, June 2009.

Communication Efficient Spatial Search in a Sensor Data Web Portal.
Yanif Ahmad, Suman Nath.
US Patent Application, February 2007.

Teaching, Mentoring and Reviewing Experience

Teaching Assistant:
CSCI0161 - Building High Performance Servers, Brown, Fall 2004, helped develop and grade homework and programming assignments for a new class on high-performance server design, including assignments emphasizing scalability aspects of simple UNIX daemons, as well as the final course project on developing an IMAP server.

Guest Lecturer:
CS0632 - Database Management Systems, Cornell, Spring and Fall 2009, lectured on centralized stream processing architectures, distributed streaming algorithms including sketching, and filtering protocols, and adaptive query processing.

Student advising:
Shubham Chopra (MSc. ORIE, Cornell, Fall 2009), Ki Suh Lee (Research Assistant, Cornell, Summer 2009, continuing as a Ph.D student in CS), Jennie Rogers (Sc.M., Brown, 2007, continuing as a Ph.D. student), Alexander Zgolinski (Sc.M. Brown, 2005, first employment Oracle Corporation), Siu Kee Kate Ho (Sc.B. Honors thesis, Brown, 2003, first employment Microsoft), George Kong (Sc.M., Brown, 2002, first employment NetApp).

External reviewer:
TKDE, VLDB Journal, PEPM 2010, VLDB 2009, SIGMOD 2009, SIGMOD 2008, ICDE 2008, VLDB 2007, ICDCS '06, ICDE '06, SIGMOD '04, ICDE '04, ICDCS '03, CIKM '02.

Student volunteer: ICDE '04, ICDCS '03.

Talks

November 2009, Oxford University, *"DBToaster: Compiling Database Systems and Queries"*.
July 2008, IBM Almaden, *"The Aurora and Borealis Stream Processing Engines"*.
June 2008, Yahoo! Research, *"Query Optimization for Efficient Sensor Stream Processing"*.
April 2008, Cornell University, *"Query Optimization for Efficient Sensor Stream Processing"*.
June 2006, Microsoft Research, *"The Borealis Distributed Stream Processing Engine"*.
2002-current, conference talks and demos, VLDB 09, ICDE 08 (Pulse and COLR-Tree), SIGMOD 06 (XPORT demo), GSN 06, SIGMOD 05, NetDB 05, VLDB 04, PODC 02.

Departmental Service

2005 – 2006, Faculty-Graduate Liaison, Brown CS. Addressed departmental issues arising concerning faculty and graduate students, including individual concerns, building and ergonomic issues, and generally acted as a voice for the graduate students.
2005 – 2006, Gong Show Czar, Brown CS. Established a volunteer position to organize events to promote graduate student research in the department, for example assisting the chair in organizing the first Departmental Retreat.
2003 – 2004, Faculty Search Czar, Brown CS. Assisted the faculty in hiring new members, and overhauled the graduate student involvement in the process, including establishing a subcommittee of 12 graduate students to aid in determining a shortlist of candidates, and in writing reviews of candidates.

References

Uğur Çetintemel
Associate Professor
Dept. of Computer Science
Brown University
Providence, RI, USA
Email: ugur@cs.brown.edu
Tel: 401-863-7644

Christoph Koch
Associate Professor
Dept. of Computer Science
Cornell University
Ithaca, NY, USA
Email: koch@cs.cornell.edu
Tel: 607-255-4117

Stan Zdonik
Professor
Dept. of Computer Science
Brown University
Providence, RI, USA
Email: sbz@cs.brown.edu
Tel: 401-863-7648

Johannes Gehrke
Professor
Dept. of Computer Science
Cornell University
Ithaca, NY, USA
Email: johannes@cs.cornell.edu
Tel: 607-255-1045

Suman Nath
Researcher
Networked Embedded Computing Group
Microsoft Research
Redmond, WA, USA
Email: sumann@microsoft.com
Tel: 425-706-8072
