

Some Query Processing Research Topics

- What are (Future) Queries and what is (Future) Query Processing?
- When/how will queries get optimized?
- How should Query Processing relate to scheduling?
- Heterogeneous Data and Big Data?
- Statistics? Access structures?
- How might Learning help?
- Code Generation and Query Execution
- Impact of emerging hardware directions; role of special purpose HW
- Security/Privacy/Trust and Query Processing?

Session 6: Query Processing Under New Assumptions

The Case Against Specialized Graph Analytics Engines

Jing Fan (University of Wisconsin); Adalbert Gerald Soosai Raj (University of Wisconsin); **Jignesh M. Patel** (University of Wisconsin)

FPGA-Based Multithreading For In-Memory Hash Joins

Robert J. Halstead (University of California); Ildar Absalyamov (University of California); Walid A. Najjar (University of California); Vassilis J. Tsotras (University of California)

Raising Authorization Awareness in a DBMS

Abhijeet Mohapatra (Stanford University); **Ravi Ramamurthy** (Microsoft Research); Raghav Kaushik (Microsoft Research)

Tupleware: "Big" Data, Big Analytics, Small Clusters

Andrew Crotty (Brown University); Alex Galakatos (Brown University); Kayhan Dursun (Brown University); **Tim Kraska** (Brown University); Ugur Cetintemel (Brown University); Stan Zdonik (Brown University)