

List of Seminar Topics

Last update: Apr 26, 2026

- **General:** Papers in the fields of **systems** for **data engineering**, **data management**, and **machine learning**
- **This Semester's Umbrella Topic: "Robust and Adaptive Query Processing"**
 - **Motivation**
 - Traditionally, **database query processing** is divided into an **optimization phase**, which determines an optimal plan for the query, and an **execution phase**, which executes this plan
 - During optimization, different **logically equivalent plans** are enumerated and the plan with the **lowest cost** with respect to some cost model is chosen
 - **Cost estimation** is largely **based on estimates of the cardinalities** (sizes) of intermediate results
 - Unfortunately, these **estimates are often quite wrong** resulting in **bad query execution plans** that may take orders of magnitude longer to execute than the optimal plan
 - Moreover, **additional unknowns** further complicate the efficient query processing, e.g., unknown properties of base data and input datasets, the access to external data sources, query parameters, and the system utilization at run-time

- **General:** Papers in the fields of **systems** for **data engineering**, **data management**, and **machine learning**
- **This Semester's Umbrella Topic: "Robust and Adaptive Query Processing"**
 - **How can we ensure efficient query processing in the presence of so many unknowns?**
 - **Ideas to address these challenges**
 - **Improved creation/management of statistics:** Try to make more accurate information on the base data and intermediate results available during query optimization
 - **Robust query optimization:** Rather than searching for an optimal query execution plan, search for a robust plan that is insensitive to inaccurate estimates
 - **Adaptive query processing:** Introduce a feedback loop to exploit information gathered at query run-time for adapting the query execution dynamically
 - **Active field of research**
 - After an **initial peak in the 2000s**, robust and adaptive query processing is **still an active research area today**
 - The seminar collects a **broad range of relevant papers** addressing various aspects of robust and adaptive query processing

■ Improved Creation/Management of Statistics

- 1) Stillger et al.: **LEO - DB2's LEarning Optimizer** (VLDB 2001) [[link](#)]
- 2) Chaudhuri and Narasayya: **Automating Statistics Management for Query Optimizers** (Trans Knowl Data Eng 2001) [[link](#)]
- 3) Wu et al.: **Sampling-Based Query Re-Optimization** (SIGMOD 2016) [[link](#)]
- 4) Dutt et al.: **Selectivity Estimation for Range Predicates using Lightweight Models** (PVLDB 2019) [[link](#)]

▪ Robust Query Optimization

- 5) Babcock and Chaudhuri: **Towards a Robust Query Optimizer: A Principled and Practical Approach** (SIGMOD 2005) [[link](#)]
- 6) Babu et al.: **Proactive Re-optimization** (SIGMOD 2005) [[link](#)]
- 7) Doraiswamy et al.: **On the Production of Anorexic Plan Diagrams** (VLDB 2007) [[link](#)]
- 8) Bizarro et al.: **Progressive Parametric Query Optimization** (Trans Knowl Data Eng 2009) [[link](#)]
- 9) Dutt and Haritsa: **Plan Bouquets: Query Processing without Selectivity Estimation** (SIGMOD 2014) [[link](#)]
- 10) Wei and Trummer: **ROME: Robust Query Optimization via Parallel Multi-Plan Execution** (SIGMOD 2024) [[link](#)]

▪ Adaptive Query Processing 1990s/2000s

11) Kabra and DeWitt: **Efficient Mid-Query Re-Optimization of Sub-Optimal Query Execution Plans** (SIGMOD 1998) [[link](#)]

12) Urhan et al.: **Cost Based Query Scrambling for Initial Delays** (SIGMOD 1998) [[link](#)]

13) Avnur and Hellerstein: **Eddies: Continuously Adaptive Query Processing** (SIGMOD 2000) [[link](#)]

14) Markl et al.: **Robust Query Processing through Progressive Optimization** (SIGMOD 2004) [[link](#)]

15) Ives et al.: **Adapting to Source Properties in Processing Data Integration Queries** (SIGMOD 2004) [[link](#)]

16) Babu et al.: **Adaptive Ordering of Pipelined Stream Filters** (SIGMOD 2004) [[link](#)]

17) Li et al.: **Adaptively Reordering Joins during Query Execution** (ICDE 2007) [[link](#)]

■ Adaptive Query Processing 2010s/2020s

- 18) Raducanu et al.: **Micro Adaptivity in Vectorwise** (SIGMOD 2013) [[link](#)]
- 19) Karnagel et al.: **Adaptive Work Placement for Query Processing on Heterogeneous Computing Resources** (PVLDB 2017) [[link](#)]
- 20) Wolf et al.: **On the Calculation of Optimality Ranges for Relational Query Execution Plans** (SIGMOD 2018) [[link](#)]
- 21) Kissinger et al.: **Adaptive Energy-Control for In-Memory Database Systems** (SIGMOD 2018) [[link](#)]
- 22) Kohn et al.: **Adaptive Execution of Compiled Queries** (ICDE 2018) [[link](#)]
- 23) Trummer et al.: **SkinnerDB: Regret-Bounded Query Evaluation via Reinforcement Learning** (SIGMOD 2019) [[link](#)]
- 24) Menon et al.: **Permutable Compiled Queries: Dynamically Adapting Compiled Queries without Recompiling** (PVLDB 2020) [[link](#)]
- 25) Justen et al.: **POLAR: Adaptive and Non-invasive Join Order Selection via Plans of Least Resistance** (PVLDB 2024) [[link](#)]