18th November 2024

QuCUN Status Seminar

Munich

QC4DB: Accelerating Relational Database Management Systems via Quantum Computing

Professor Dr. rer. nat. habil. Sven Groppe
https://www.ifis.uni-luebeck.de/index.php?id=groppe

QC4DB: Accelerating Relational Database Management Systems via Quantum Computing

Name:	QC4DB: Accelerating Relational Database Management Systems via Quantum Computing	
Proj. Web:	**Project Website@Quantentechnologien**	
Funded by:	BMBF, Fördermaßnahme Anwendungsnetzwerk für das Quantencomputing	
Duration:	3 years, 1.1.2022 - 30.6.2025 (extended)	
Volume:	1.8M Euros	
Topics:	Optimizing an open source relational database management system • Queries • Transaction Schedules	
Partners:	UNIVERSITÄT ZU LÜBECK INSTITUT FÜR INFORMATIONSSYSTEME (Coord.)	QUANTUM BRILLIANCE
Expertises:	Hardware-Acceleration of Databases	Room Temperature Diamond Quantum Accelerators/qbOS
Website:	**https://www.ifis.uni-luebeck.de/~groppe/**	**https://quantumbrilliance.com/**

Timeline of Activities

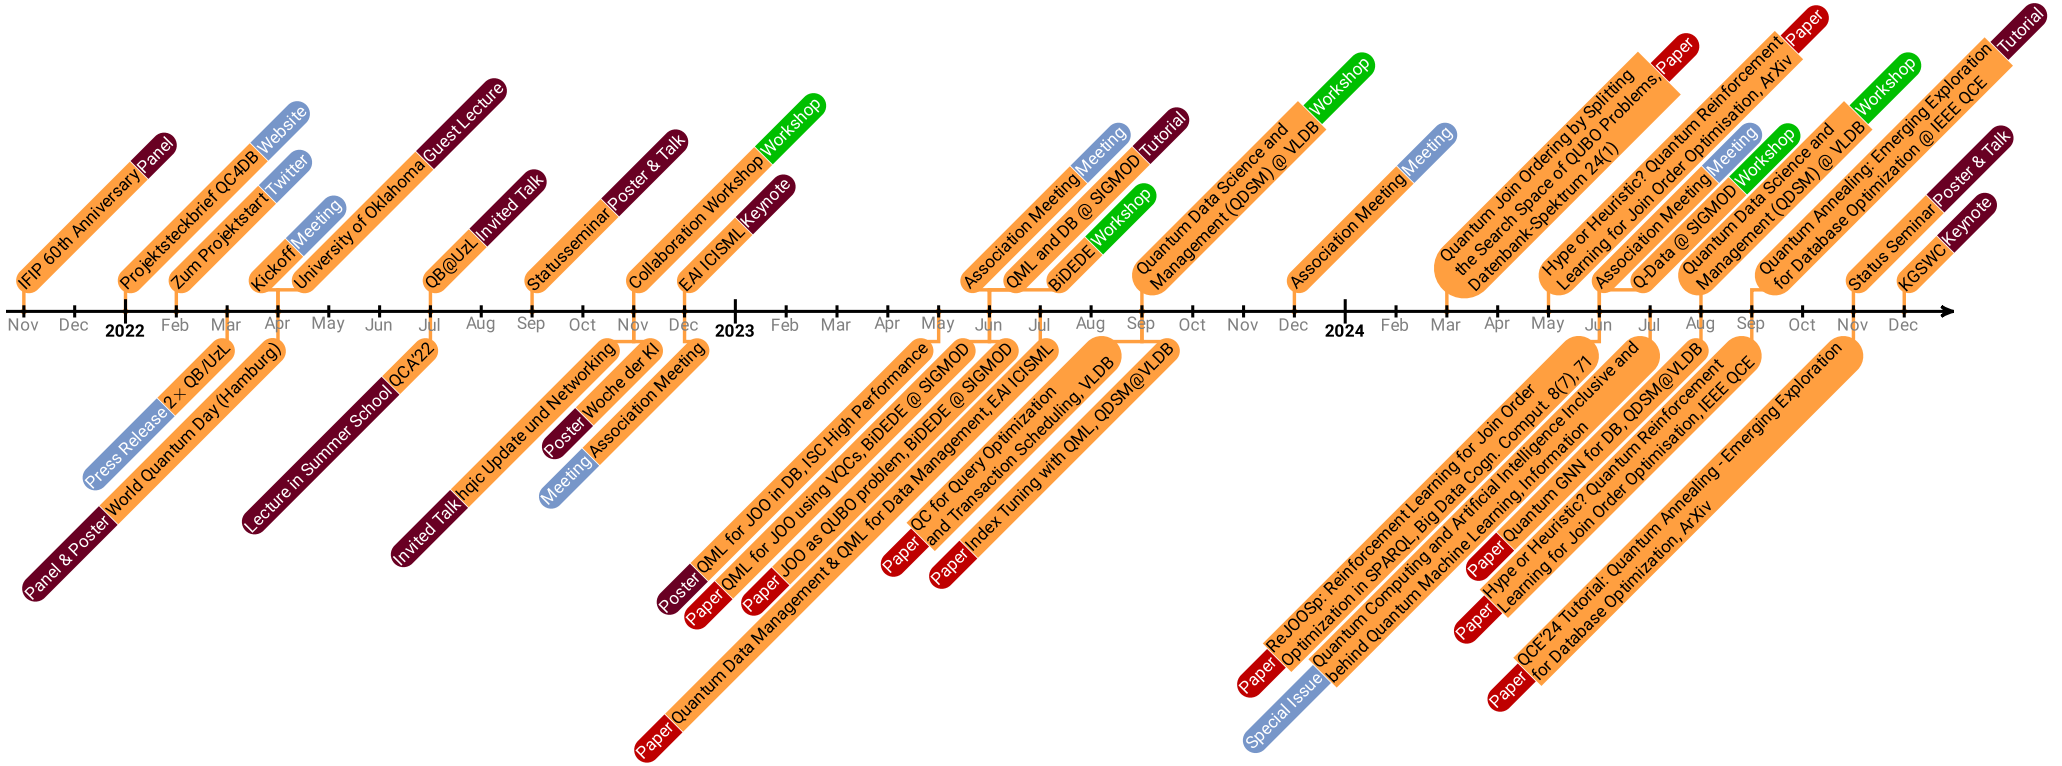

- 11 Papers
- 5 Workshops
- 2 Tutorials
- 5 Posters
- 2 Panels
- 2 Keynotes
- 5 Invited Talks
- 1 Special Issue in Journals

Number of Citations (18.11.2024)

#Citations	Publication
28	Quantum machine learning: Foundation, new techniques, and opportunities for database research
24	Constructing optimal bushy join trees by solving qubo problems on quantum hardware and simulators
21	Opportunities for quantum acceleration of databases: optimization of queries and transaction schedules
12	Quantum machine learning for join order optimization using variational quantum circuits
15	Index Tuning with Machine Learning on Quantum Computers for Large-Scale Database Applications
6	Quantum data management and quantum machine learning for data management: State-of-the-art and open challenges
1	Hype or Heuristic? Quantum Reinforcement Learning for Join Order Optimisation
1	Quantum Join Ordering by Splitting the Search Space of QUBO Problems
0	ReJOOSp: Reinforcement Learning for Join Order Optimization in SPARQL
2	Supervised Learning on Relational Databases with Quantum Graph Neural Networks
0	QCE'24 Tutorial: Quantum Annealing - Emerging Exploration for Database Optimization
$\sum$: 110	in total

according to Google Scholar

Importance of Collaboration

- Collaboration Workshop in Nov. 2022, Participants: Groups of
  1. Prof. Jiaheng Lu (University of Helsinki, Finland)
  2. Prof. Wolfgang Mauerer (OTH Regensburg, Germany)
  3. Prof. Le Gruenwald (University of Oklahoma, USA)
  4. Dr. Florian Preis (Quantum Brilliance, Germany)
  5. Prof. S. Groppe (University of Lübeck, Germany)
  6. ...
- Collaboration Activities ("result" of the collaboration workshop)
  - 4 Workshops: QDSM@VLDB '23/'24 (Org.: 1., 2., 3., 5.), BiDEDE@SIGMOD'23 (Org.: 3., 5.), QData@SIGMOD'24 (Org.: 4., Steering Committee: 2., 5.)
  - 2 Tutorials: SIGMOD'23 (1., 2., 5.), IEEE QCE'24 (1., 2., 5.)
  - 6 Papers: IEEE QCE'24 (2., 5.), SIGMOD'23 (1., 2., 5.), QDSM@VLDB'23 (3., 5.), BiDEDE@SIGMOD'23 (3., 5.), VLDB'23 (3., 4., 5.), EAI ICISML'22 (3., 5.)
- Other Collaborations:
  - 2 Papers: QDSM@VLDB'24 (5., Martin Vogrin, University of Maribor, Slovenia), BDCC 8(7) (5., Shridevi Krishnakumar, VIT Chennai, India)

Using Hardware Accelerator for optimizing Queries / Transaction Schedules

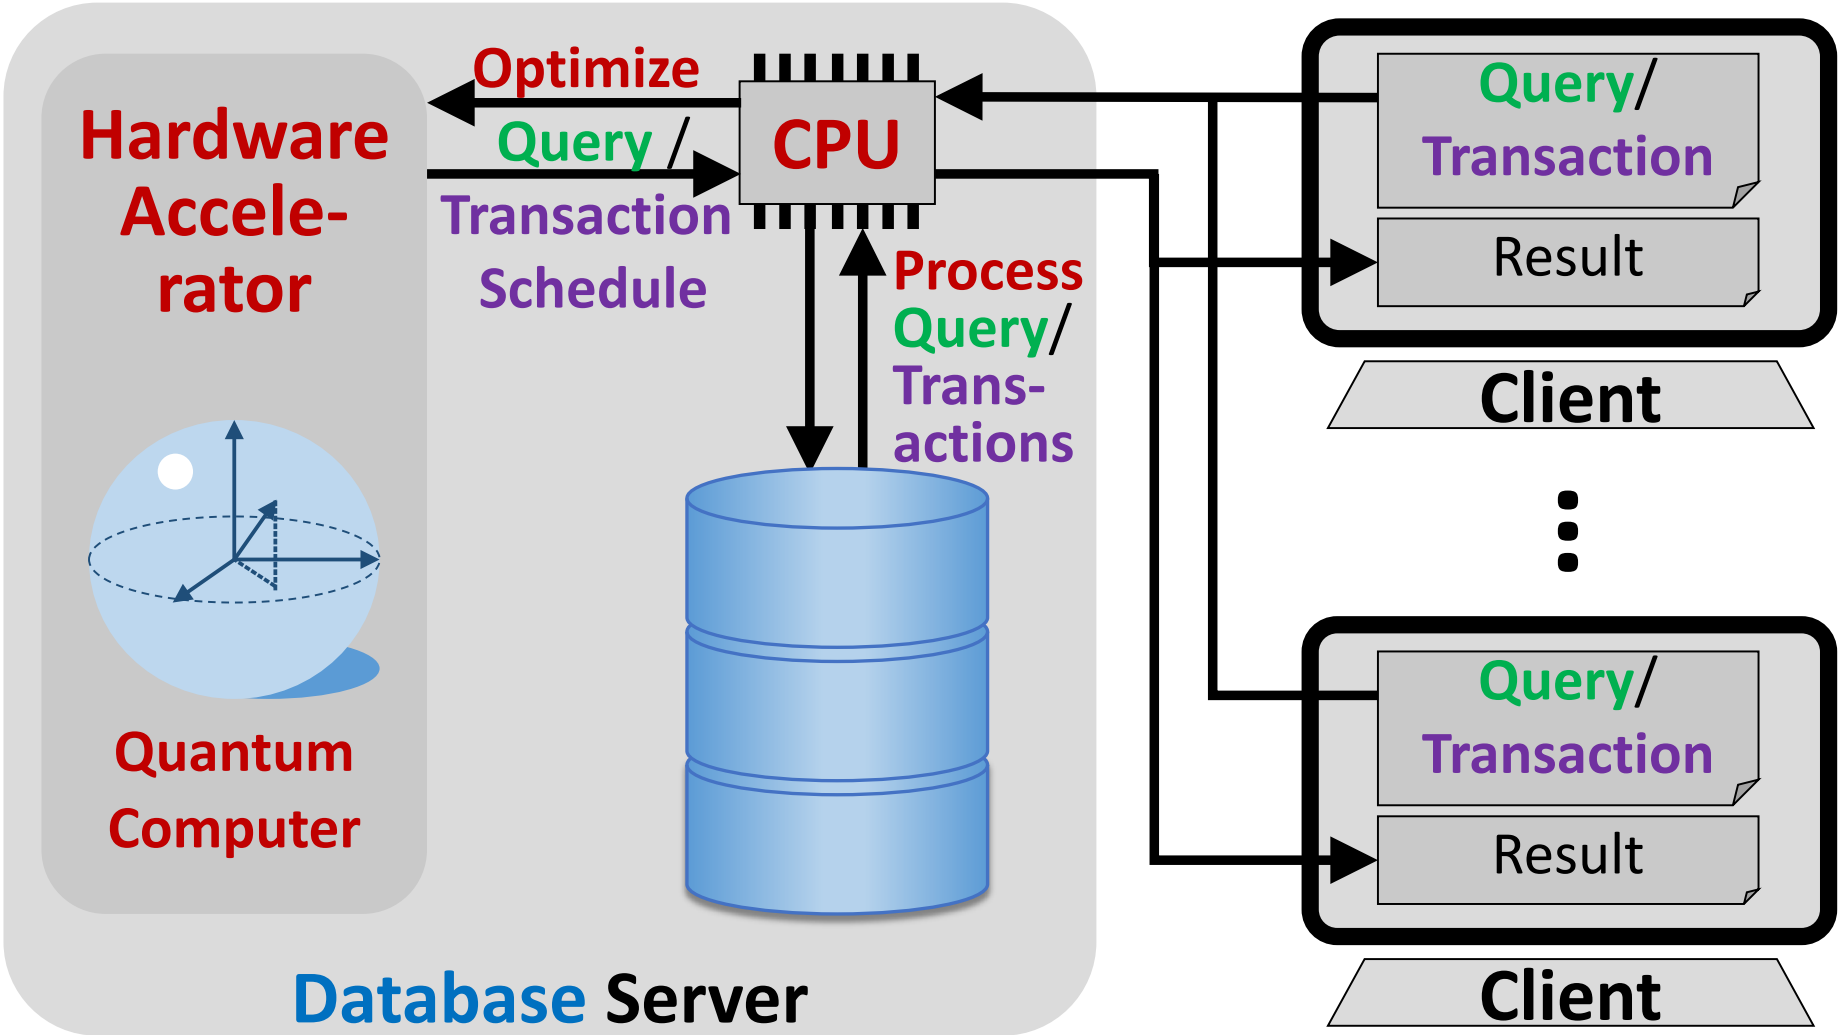

Planned Contributions

Query Optimization:

$$\bowtie_{i=1}^{n} R_i \rightarrow ?$$

$$(R_1 \bowtie R_2) \cdots \bowtie R_n$$
$$\vdots$$
$$(R_1 \bowtie R_n) \bowtie (\cdots)$$

Transaction Schedule Optimization:

$$\{T_1, \ldots, T_m\} \rightarrow ?$$

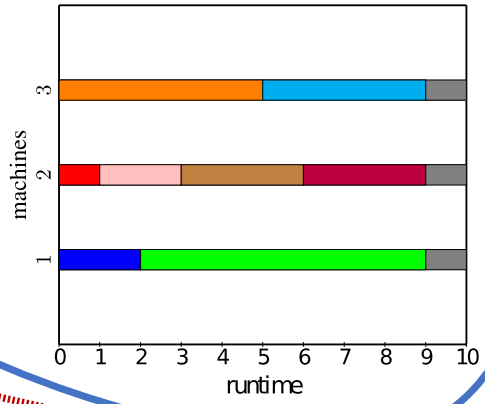

Open Source Relational Database Management System (RDBMS), e.g. PostgreSQL, MySQL

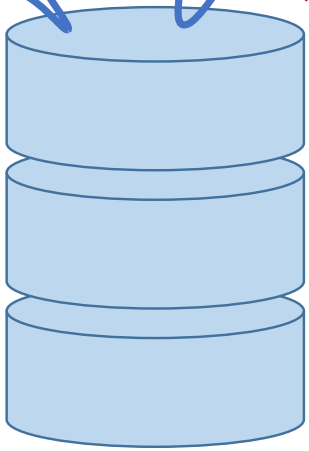

- Dynamic Programming
- Random Walk
- Linear Program-ming
- ...
- Machine Learning
- Genetic Algorithm
- Simulated Annealing

Transaction Scheduling

- Design choice: Fast quantum exact methods and good heuristics for scalable optimization times
- Classical
  - Dynamic Programming | Random Search | Simulated Annealing (SA) | Genetic Algorithm (GA) | Ant Colony Optimization (ACO)
- Quantum
  - Maximum Independent Set | Quantum Evolutionary Algorithms (QEA) (several variants: Binary | (Cultural (QCEA)) Multi-State) | QUBO | Grover

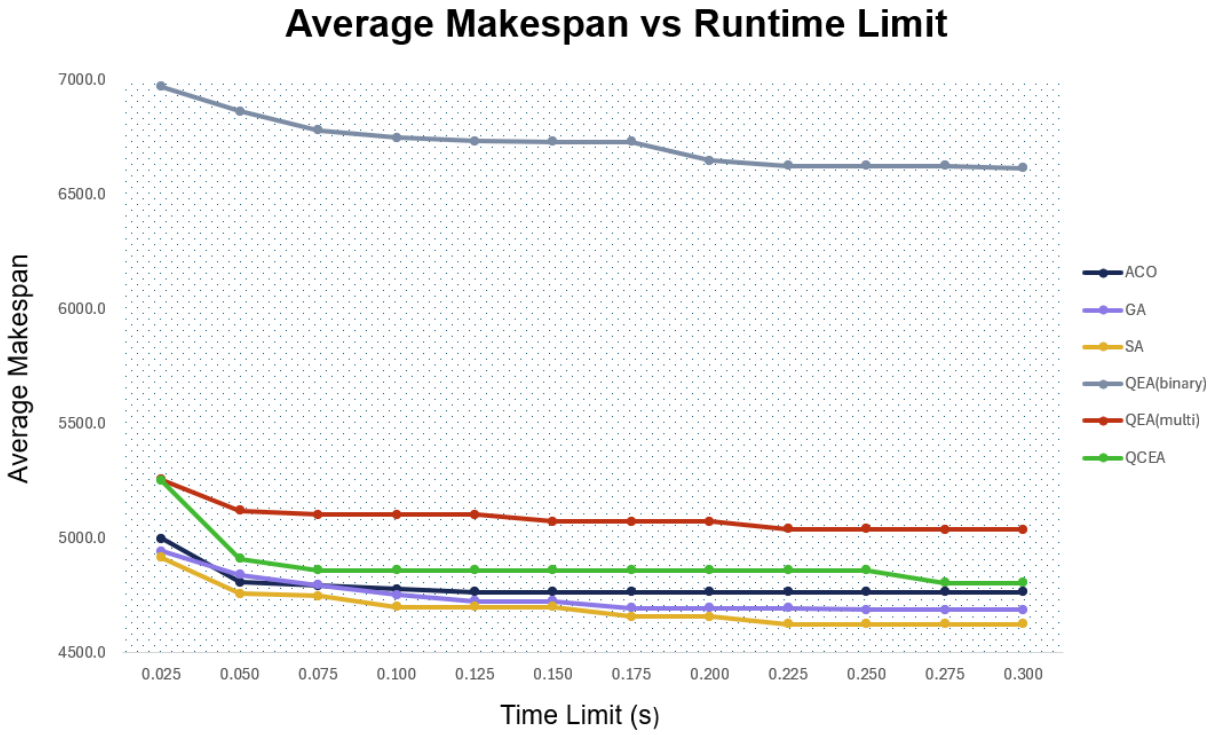

Average values for 100 problem instances with 30 transactions, 4 cores and 0.5 conflict probability

QML for Join Ordering

- Design choice: Predict best join order based on real execution times in contrast to estimated costs (EC)
  - can beat exact methods like dynamic programming (DP) based on EC
- Single-Step: Real-world queries (ErgastF1 Benchmark/PostgreSQL)
  - join orders with faster execution than DP (2.8%) and GEQO (11.2%)
  - close to classical ML like RTOS (6%) / best join orders (16.8%)

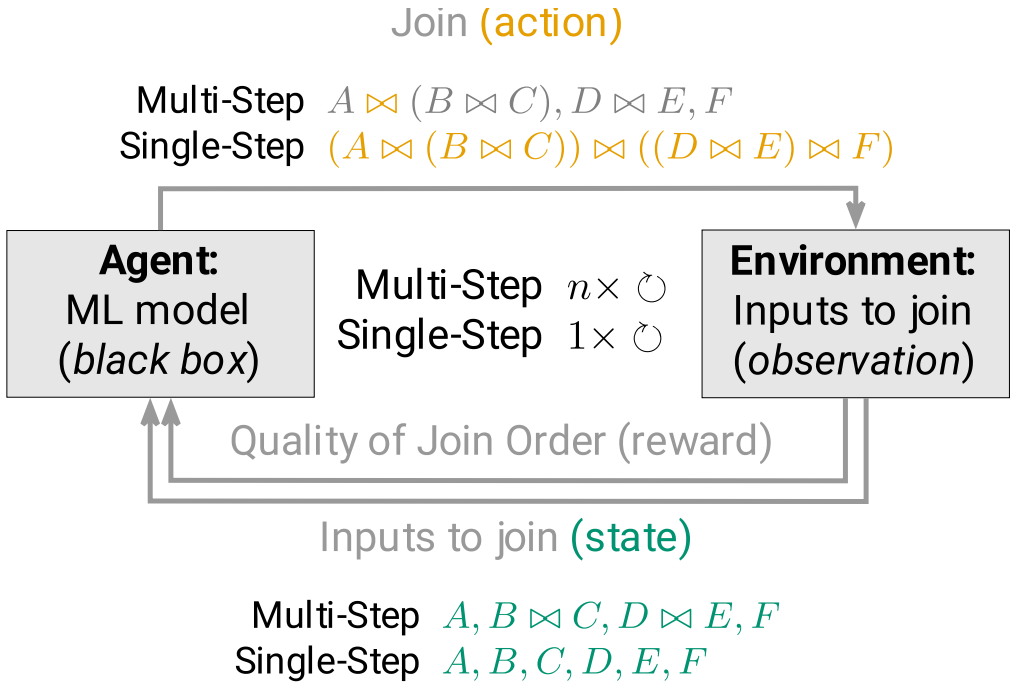

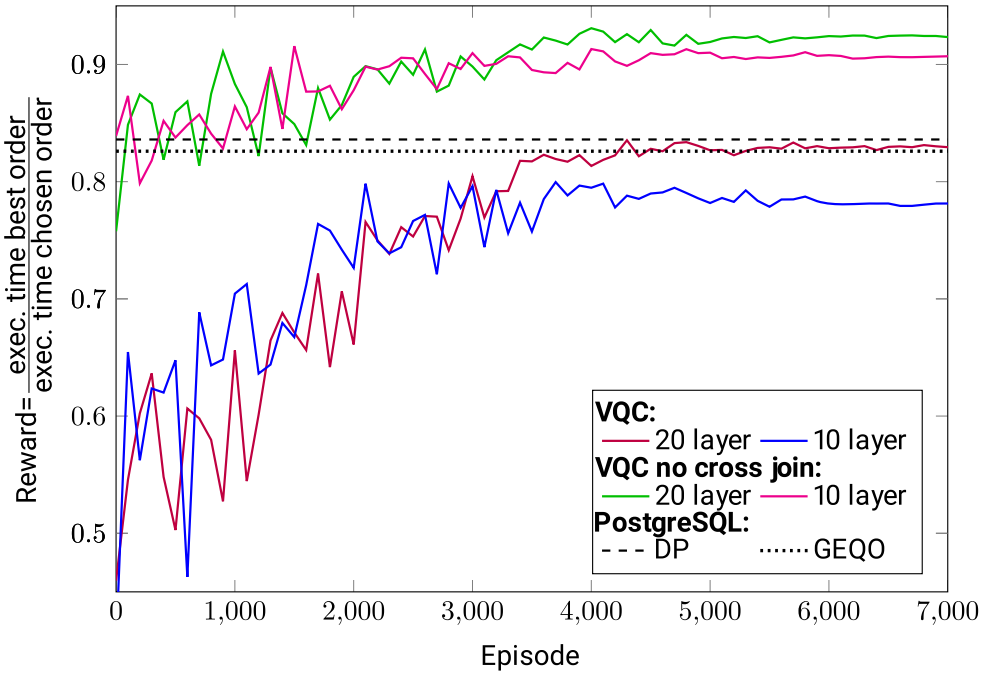

Join Ordering as QUBO Problem

- Design choice: Direct Join Costs in contrast to calculated costs
  - Estimated costs for **each** join $\Rightarrow$ better join order
  - Runtime complexity $O(2^m - m)$ with m relations to join optimal for join ordering based on direct costs
    - Classical runtime: $O(3^m - 2^{m+1})$ using dynamic programming
- Real-world queries (ErgastF1 Benchmark/PostgreSQL)

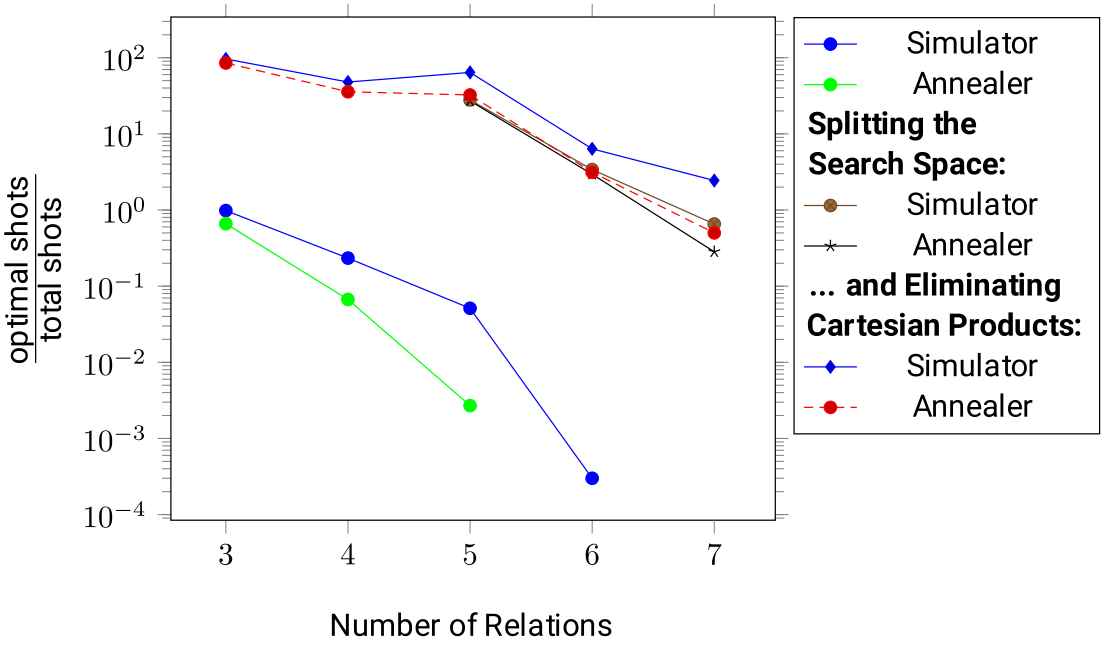

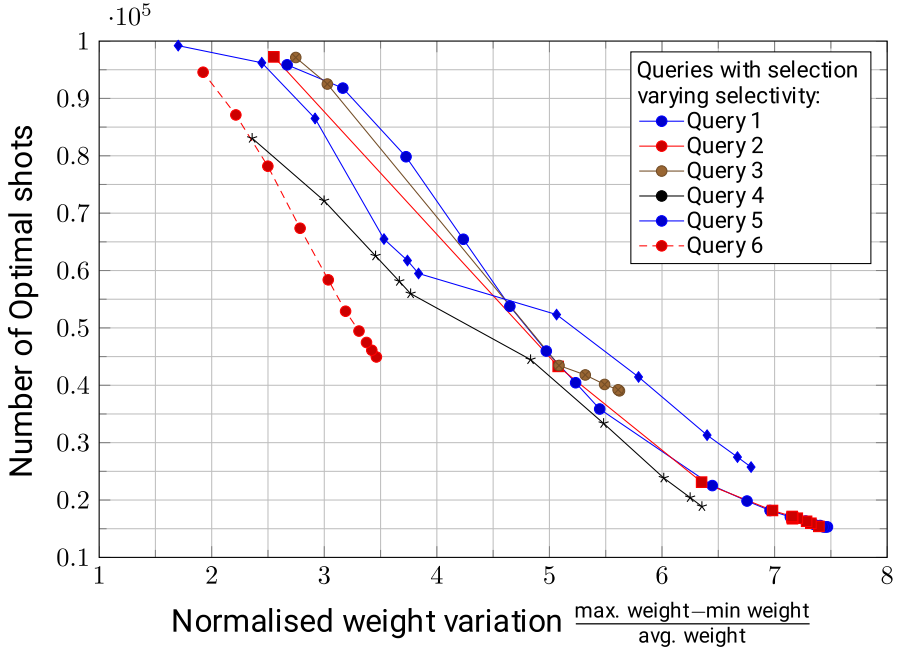

QML for Cardinality Estimation

- Design choice: Predict correction factor of a Cardinality Estimator
  - Cardinality Predictor:PostgreSQL|MSCN|Hybrid Quantum Classical Network
  - Cardinality Correction: Hybrid Quantum Classical Network
- Queries of JOB-light Benchmark

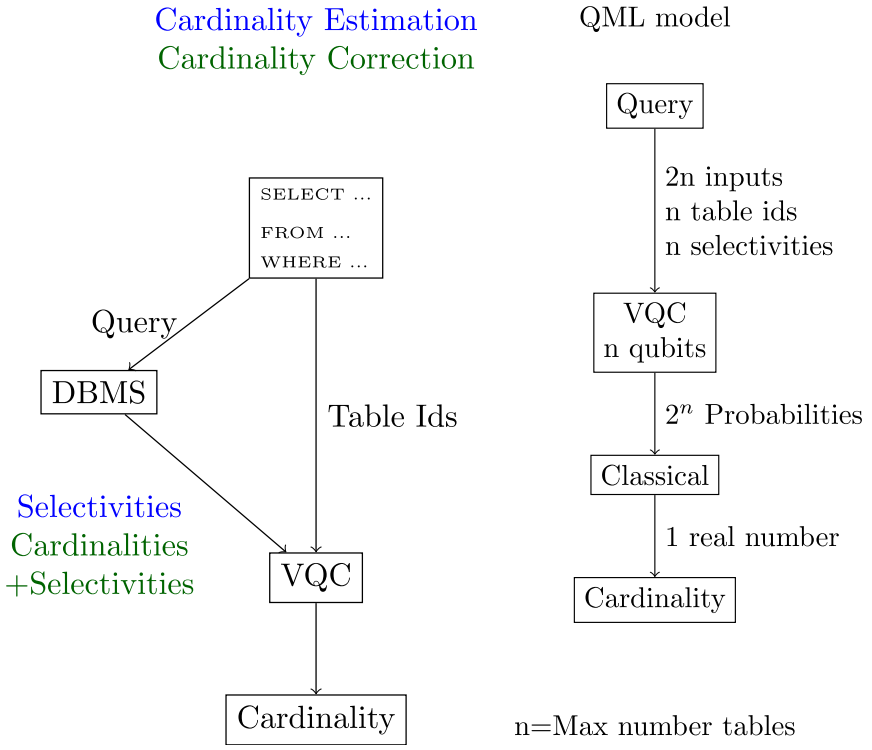

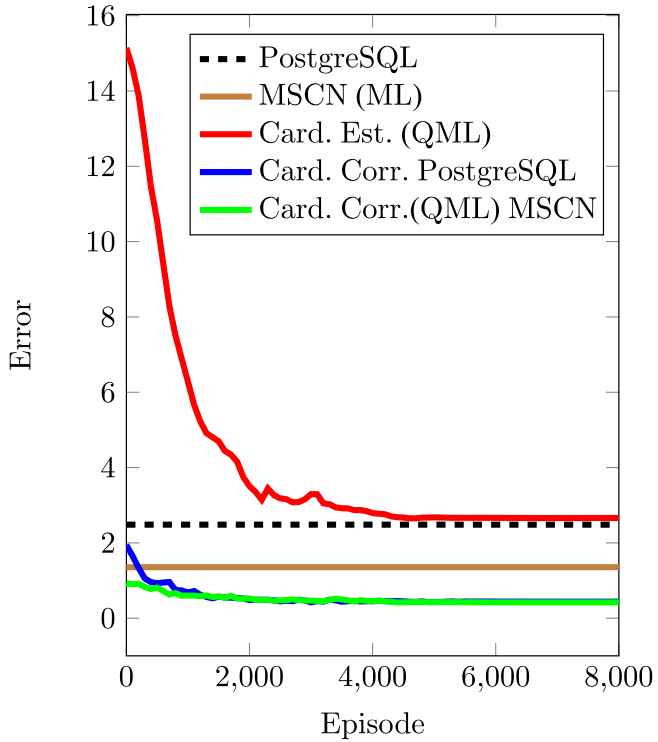

Transaction Scheduling Optimisation Server

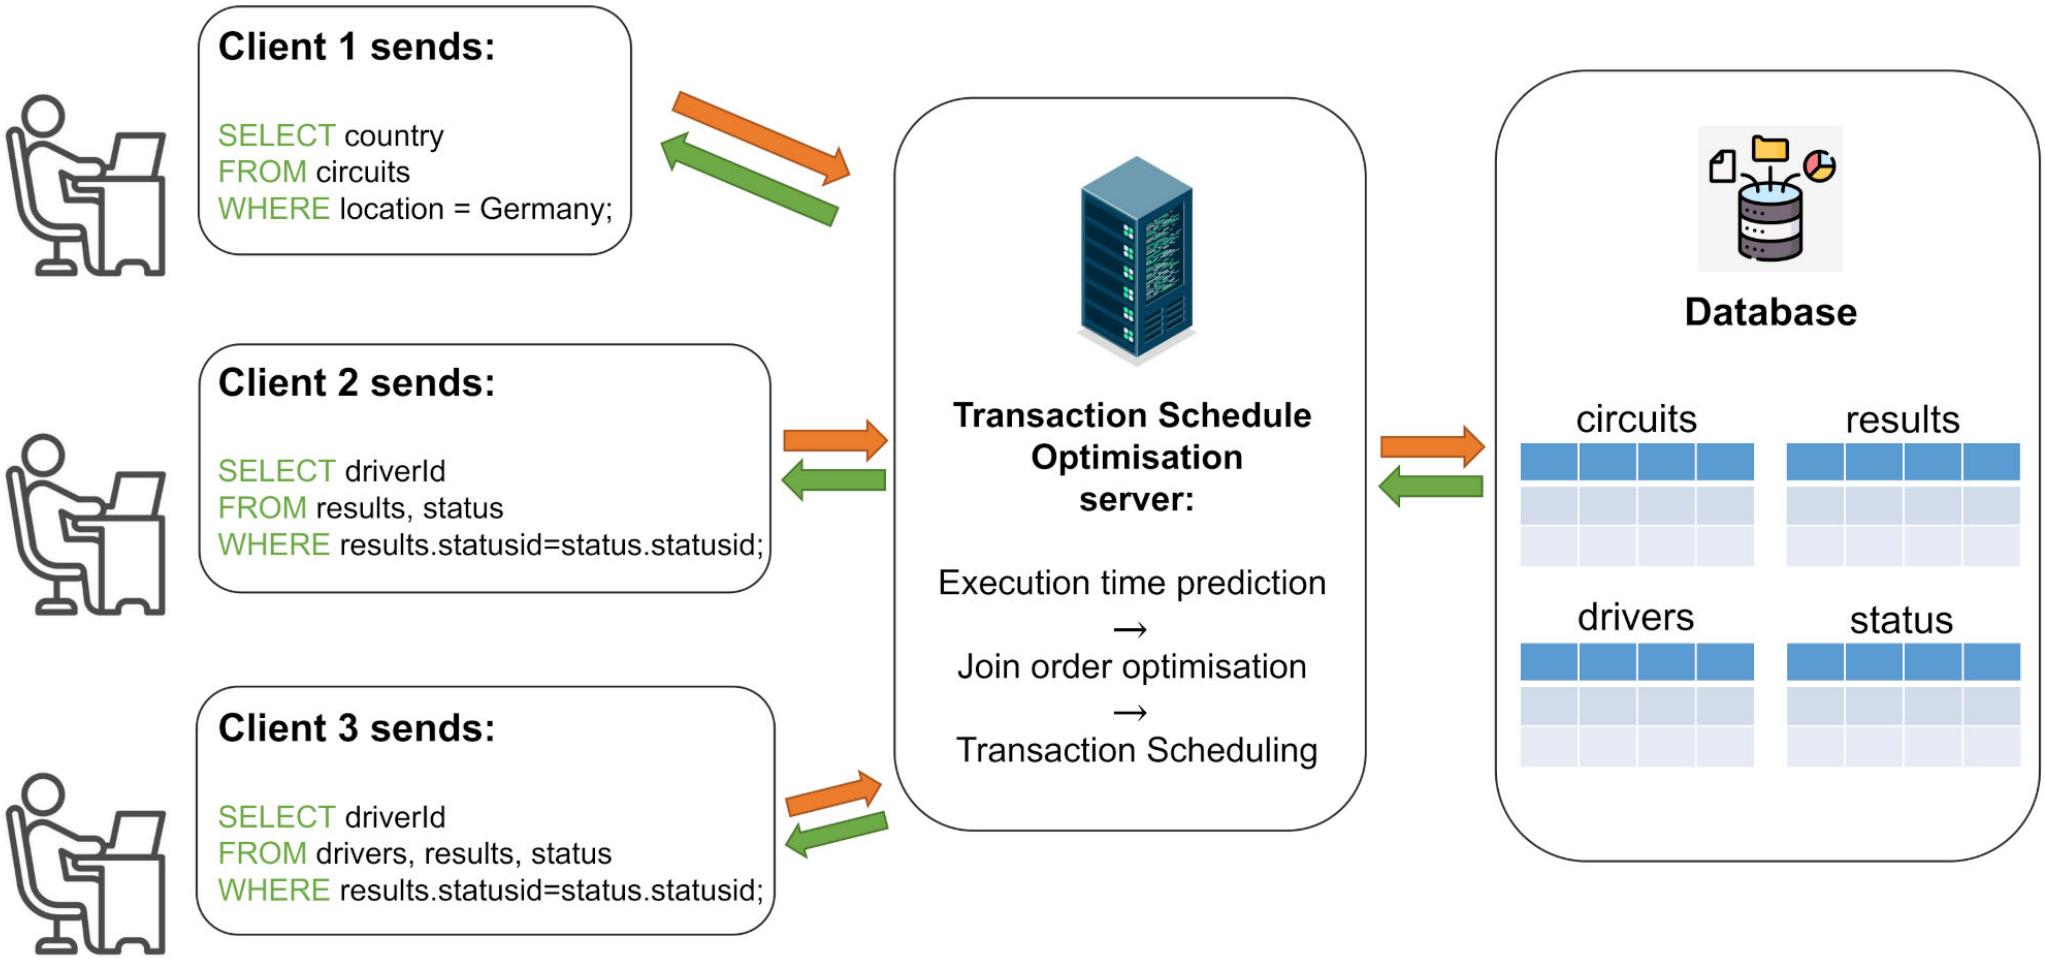

- Demonstrator
  - fully quantum **versus** fully classical

Demonstrator

- is publicly available at the Quantum Brilliance web-page...

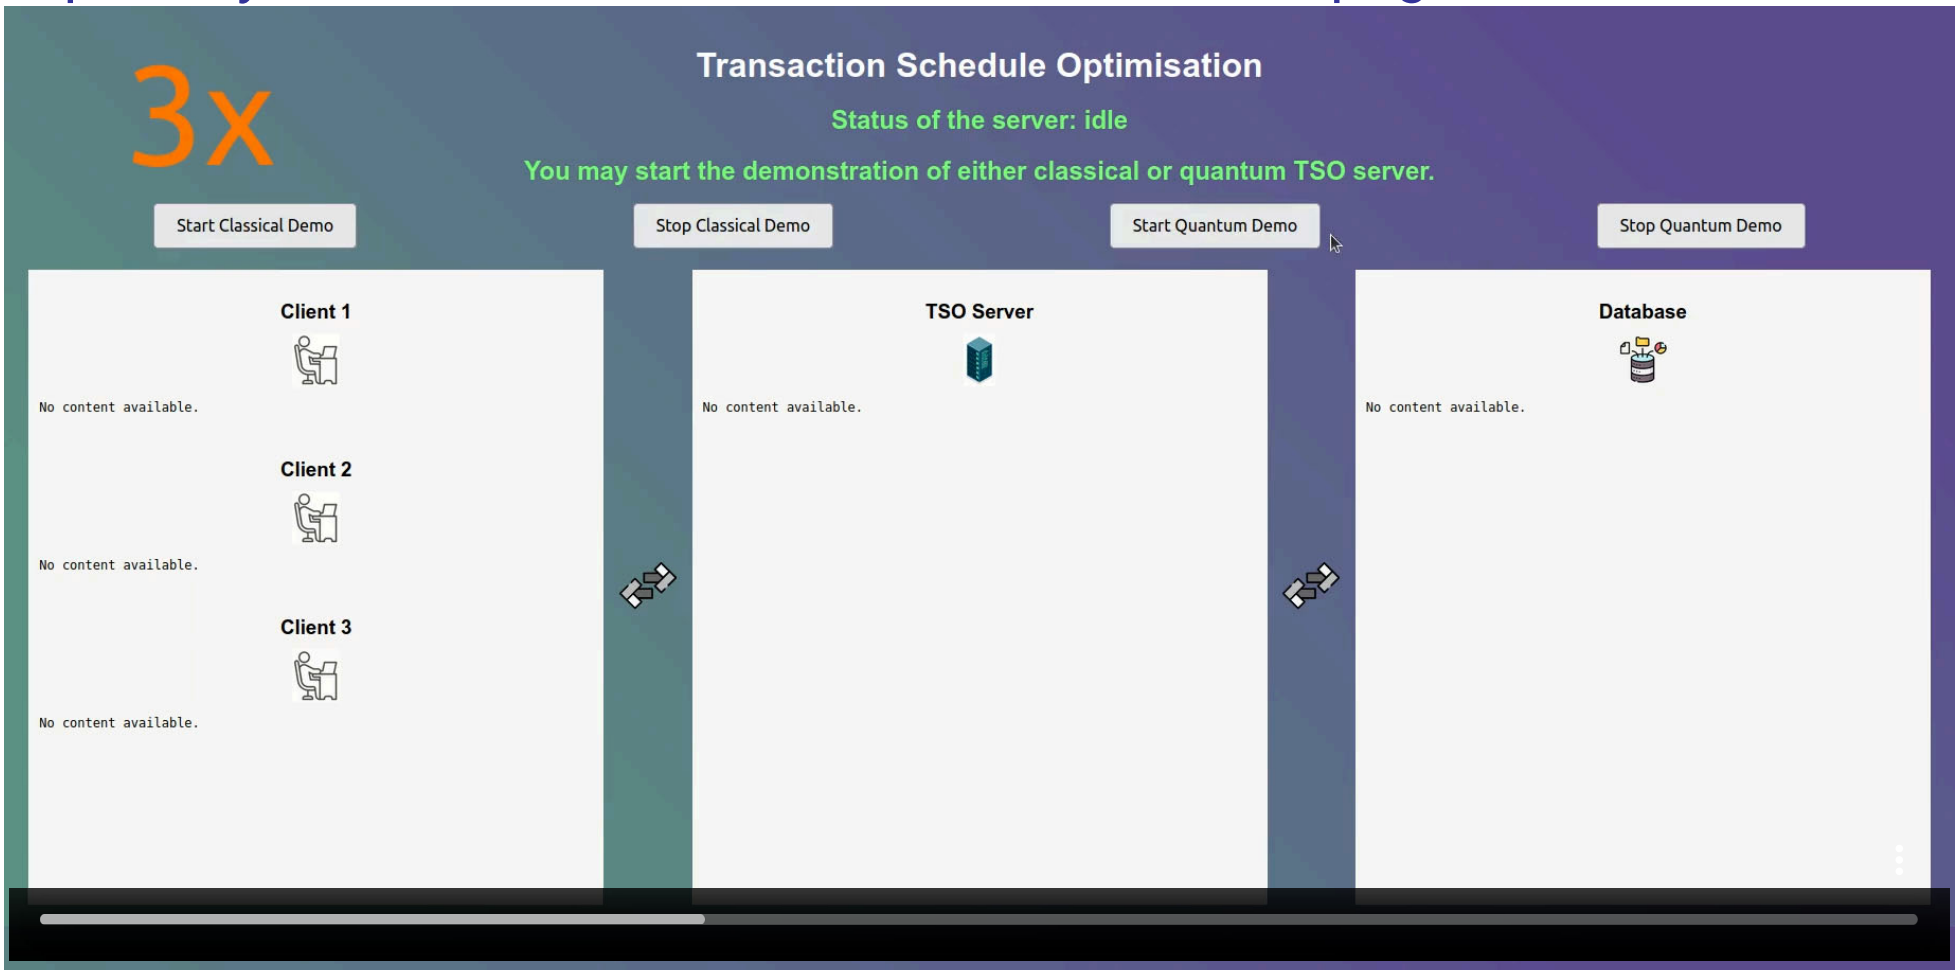

QC4DB Team

Partner:	University of Lübeck	Quantum Brilliance GmbH
Project Leaders:	Prof. Dr. rer. nat. habil. Sven Groppe (Coordinator)	Dr. Florian Preis Dr. Stefan Prestel Dr. Farida Shagieva
Project Employees:	Tobias Winker Umut Çalıkyılmaz Nitin Nayak Benjamin Warnke	Daanish Arya Dr. John Helm Dr. Johannes Kuhlmann Dr. Jeremy Rodriguez