FG DB Spring Symposium 2022, Potsdam

Quantum Computing for DB

Applicability on Multi Query Optimization and Join Order Optimization

Manuel Schönberger[1] Stefanie Scherzinger[2] Wolfgang Mauerer[1,3]

[1] Technical University of Applied Sciences Regensburg [2] University of Passau [3] Siemens AG, Corporate Research

Quantum Computing

Overview

- Quantum processing units (QPU)
- Qubits not limited to the states 0 or 1
- Speedups via quantum phenomena
  - Superposition
  - Entanglement
- Applications: ML [1], Optimization, ...

Current Challenges

- Limited qubit numbers
- Limited qubit connectivity
- Limited coherence times
- Possibility of gate errors
- Efficient problem reformulations
- Experiment reproducibility [2]

QPU DBMS Architecture Integration

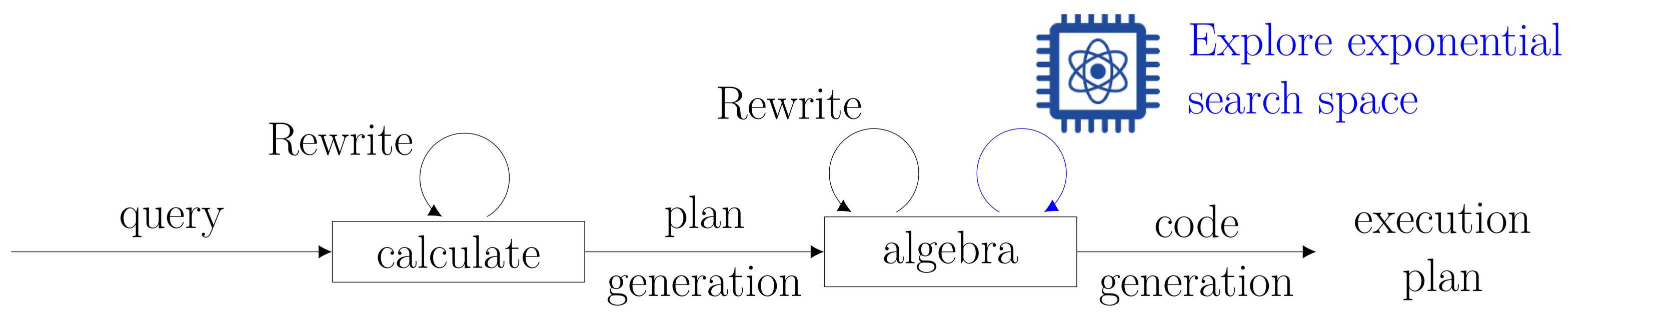

Gate-based QPUs (IBM-Q) [3]

- Up to 127 physical qubits
- Quantum circuits
- Universal computation

Quantum Annealing (D-Wave) [4]

- $\approx$ 5,000 physical qubits
- Open debate: Actual speedups over CPUs?
- Restricted to QUBO problems

Multi Query Optimization on D-Wave

VLDB'16, Trummer & Koch

Overview

- Goal: Minimize execution costs for a batch of queries
- Valid solution: One plan per query
- Naive approach: Choose locally cheapest plan
- Better: Select plans with common subexpressions

QUBO Reformulation

- Energy formula: $\omega_L E_L + \omega_M E_M + E_C + E_S$ [5]
- QUBO terms for incentivizing valid and optimal solutions:
  - $E_L = -\sum_{p \in P} X_p$
  - $E_M = \sum_{q \in Q} \sum_{\{p1,p2\} \subseteq P_q} X_{p1} X_{p2}$
  - $E_C = \sum_{p \in P} c_p X_p$
  - $E_S = -\sum_{\{p1,p2\} \subseteq P} s_{p1,\,p2} X_{p1} X_{p2}$

	Query 0			Query 1		Query 2		
Plans	p0	**p1**	p2	**p3**	p4	p5	p6	**p7**
Costs	10	**12**	15	**9**	16	7	12	**9**

Combination	**p1,p3**	**p1,p7**	p2,p3	p4,p6	p4,p7
Savings	**4**	**5**	6	7	3

Solving MQO on IBM-Q QPUs

New!

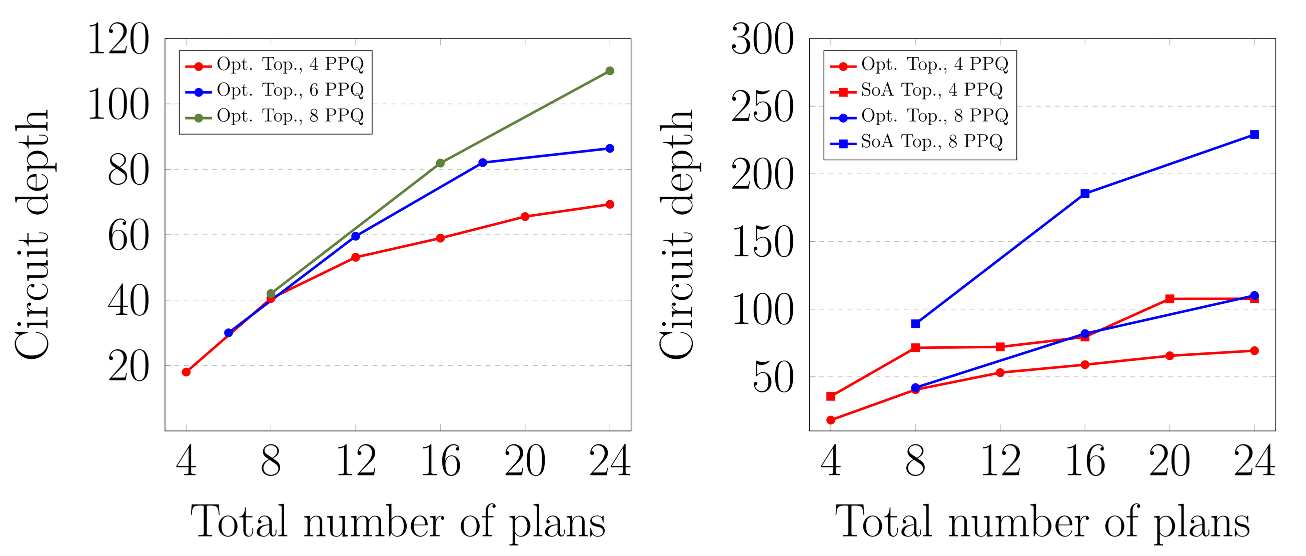

- Solving multi query optimization (MQO) with QAOA on gate-based QPUs [6, 7]
- Maximum coherent ciruit depth for the Mumbai QPU: 248 gates
- All circuits can be executed within the coherence time

Join Order Optimization

Classification
- General query graphs
- Left-deep join trees
- Support of cross products
- Minimizing intermediate cardinalities

Query graph

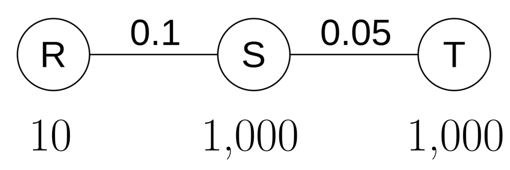

Possible join trees

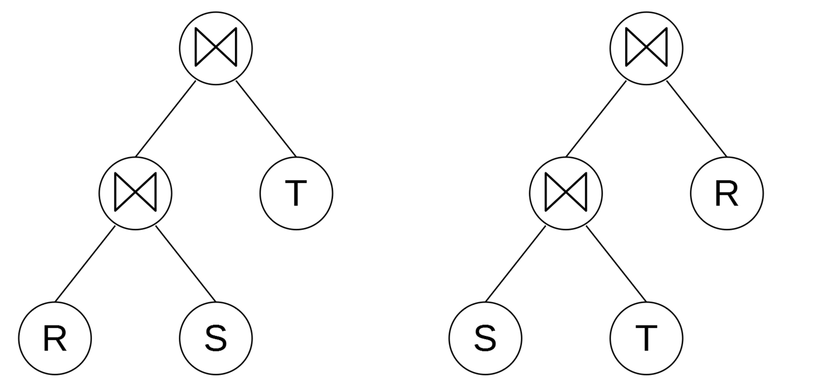

Join orders	$(R \bowtie S) \bowtie T$	$(S \bowtie T) \bowtie R$
Resulting costs	1,000	50,000

Join Ordering Reformulation

New!

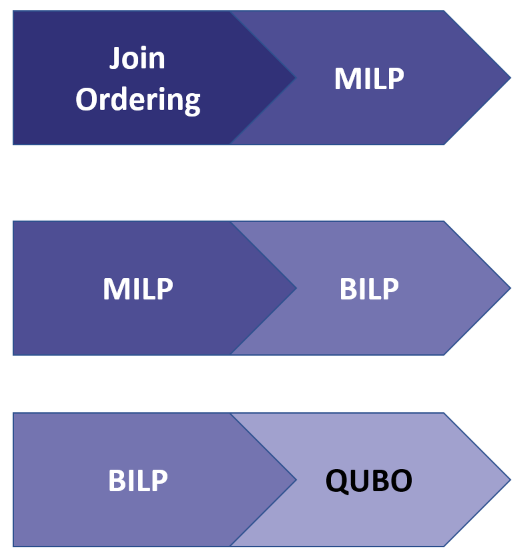

- Mixed integer linear programming (MILP) reformulation [8]
- Approximate log. cardinalities: $\min \sum_{r=0}^{R-1} \sum_{j=1}^{J-1} cto_{rj} \theta_r$
- Approx./ validity constraints: $c_j - cto_{rj} \cdot \infty_{rj} \leqslant \log(\theta_r)$, ...
- Equality conversion: $c_j - cto_{rj} \cdot \infty_{rj} + s_{rj} = \log(\theta_r)$
- Variable discretization: $s_{rj} \approx \omega \sum_{i=1}^{n} 2^{i-1} b_i$
- Transform the binary ILP (BILP) problem to QUBO [9]
- Energy formula: $A \sum_{j=1}^{m} \left(b_j - \sum_{i=1}^{N} S_{ji} x_i \right)^2 + B \sum_{i=1}^{N} c_i x_i$

Join Ordering Analyzed for QPUs

New!

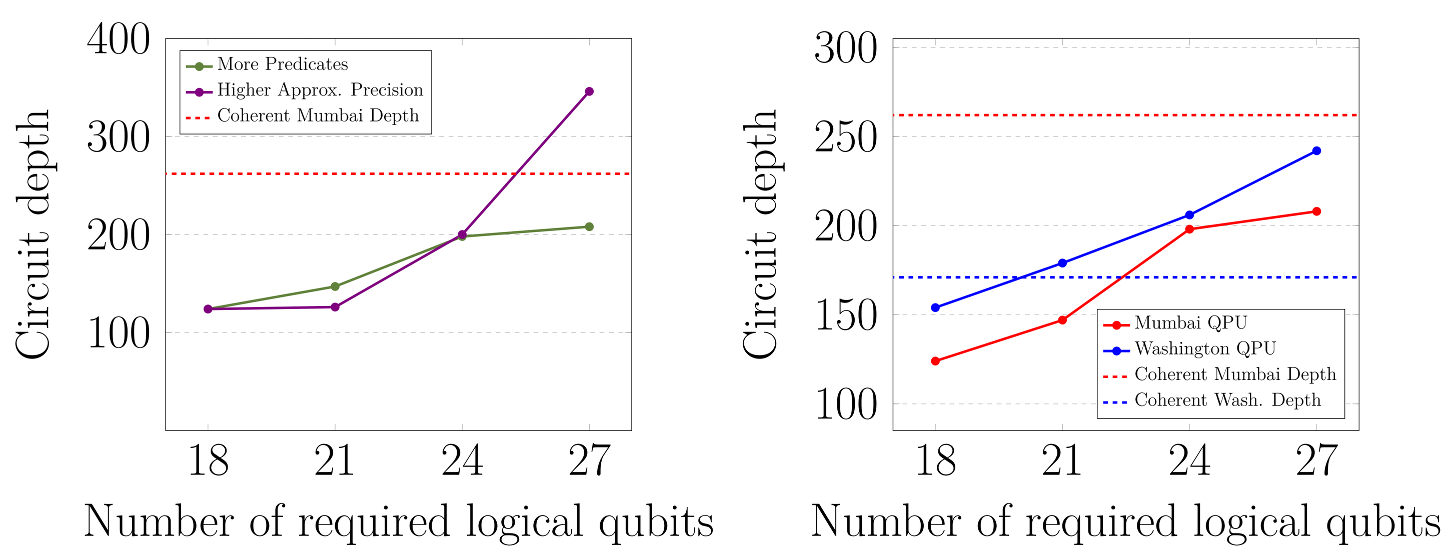

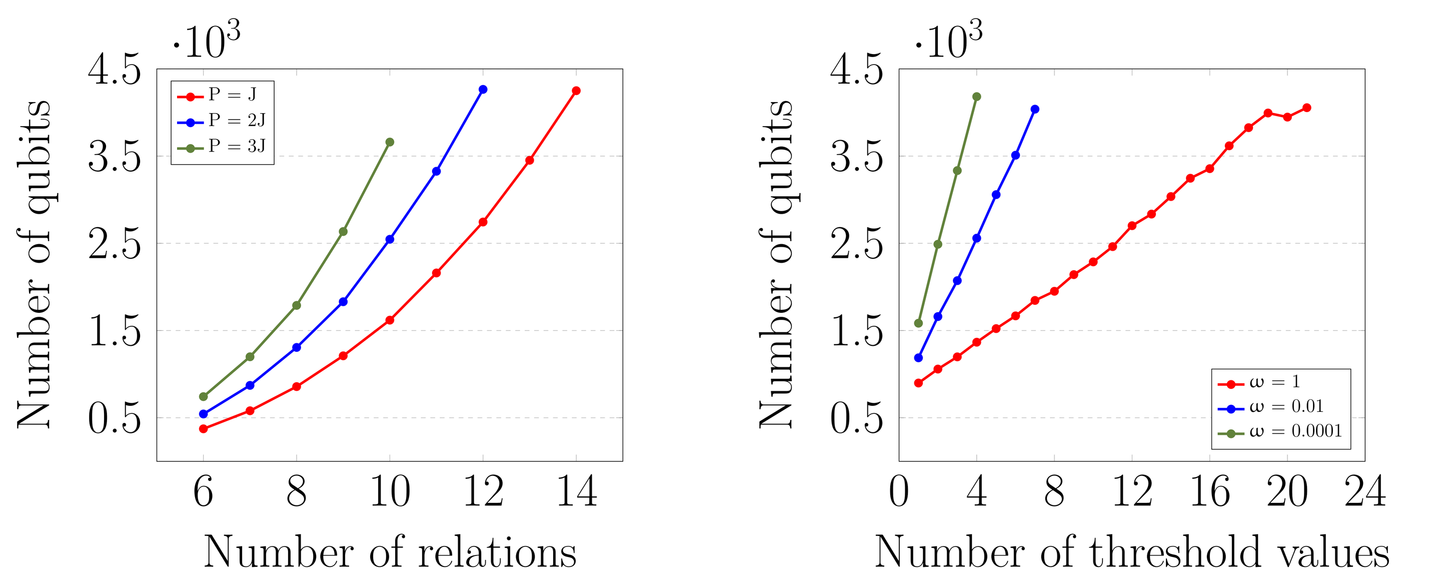

- Larger problems solvable on D-Wave systems compared to current IBM-Q QPUs
- Overall: Solvable problem dimensions still very limited
- But: Good opportunity to investigate further problems in anticipation of future QPUs

References & Funding

Own Publications

[1] Maja Franz, Lucas Wolf, Maniraman Periyasamy, Christopher Ufrecht, Daniel D. Scherer, Axel Plinge, Christopher Mutschler, and Wolfgang Mauerer. "Uncovering Instabilities in Variational-Quantum Deep Q-Networks". In: (2022). arXiv: 2202.05195 [quant-ph]. URL: https://arxiv.org/abs/2202.05195.

[2] Wolfgang Mauerer and Stefanie Scherzinger. *1-2-3 Reproducibility for Quantum Software Experiments*. 2022. arXiv: 2201.12031 [cs.SE]. URL: https://arxiv.org/abs/2201.12031.

[6] Manuel Schönberger, Maja Franz, Stefanie Scherzinger, and Wolfgang Mauerer. "Peel | pile? Cross-framework portability of quantum software". In: *19th IEEE International Conference on Software Architecture Companion (ICSA-C)*. Honolulu, HI, USA: IEEE, 2022.

[7] Manuel Schönberger. "Applicability of Quantum Computing on Database Query Optimization". In: *SIGMOD'22: International Conference on Management of Data*. Philadelphia, NY, USA: ACM, in press.

External Publications

[3] IBM Quantum. *Cloud access to quantum computers provided by IBM*. 2021. URL: https://quantum-computing.ibm.com.

[4] Catherine McGeoch and Pau Farré. *The D-Wave Advantage system: An overview*. Tech. rep. 14-1049A-A. D-Wave Systems Inc, 2020.

[5] Immanuel Trummer and Christoph Koch. "Multiple query optimization on the D-Wave 2X adiabatic quantum computer". In: *Proceedings of the VLDB Endowment* (2016).

[8] Immanuel Trummer and Christoph Koch. "Solving the join ordering problem via mixed integer linear programming". In: *Proceedings of the 2017 ACM International Conference on Management of Data*. New York, NY, USA: ACM, 2017.

[9] Andrew Lucas. "Ising formulations of many NP problems". In: *Frontiers in Physics* (2014).

This research was supported by the German Federal Ministry of Education and Research (BMBF), funding program "quantum technologies—from basic research to market", grant number 13N15645.

Federal Ministry of Education and Research