OSTBAYERISCHE TECHNISCHE HOCHSCHULE REGENSBURG

2022 ACM SIGMOD SRC, Philadelphia

Quantum Computing for DB

Applicability on Multi Query Optimization and Join Order Optimization

Manuel Schönberger[1]

[1] Technical University of Applied Sciences Regensburg

QPU DBMS Architecture Integration

Overview on Quantum Computing

- Quantum processing units (QPU) explore search spaces using quantum phenomena
- Ideal application: Query optimization, which features large search spaces
- Further applications: ML [1], Simulation

Achieving QPU-DBMS Integration

- QPU imperfections limit current utility
- To facilitate DBMS-QPU integration, we
  - show how to solve problems on QPUs
  - analyze our approach for current QPUs
  - derive criteria for DB-QPU codesign

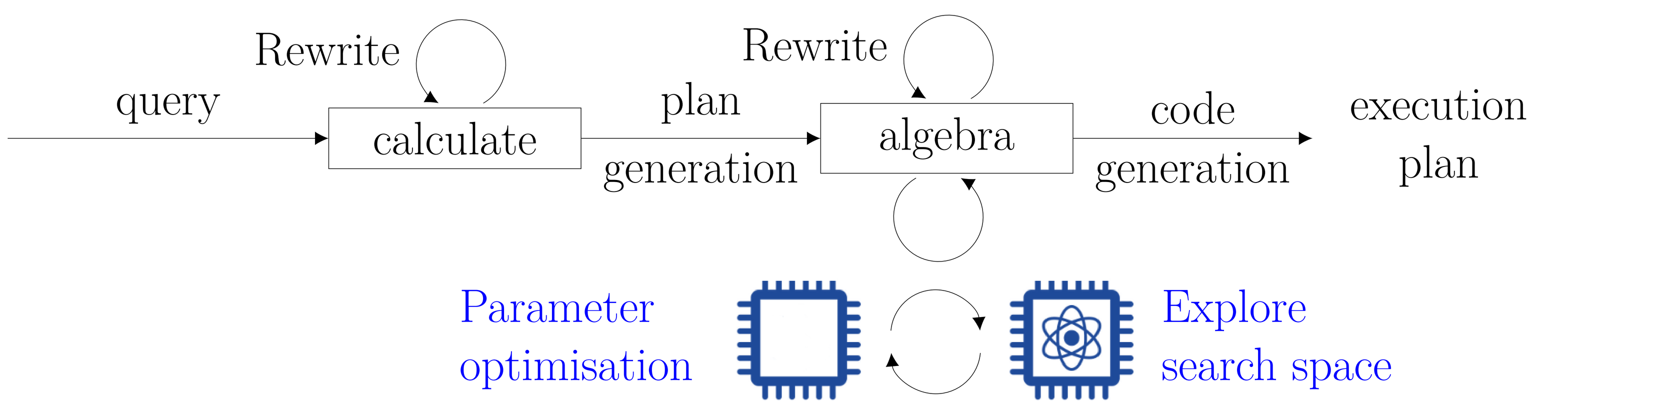

Quantum Optimization Algorithms

Quantum Approx. Opt. Algorithm (QAOA)

- Circuit for gate-based QPUs (e.g., IBM-Q)
- Hybrid quantum-classical algorithm:
  - QPU: Parameterized state preparation
  - CPU: Parameter optimization based on $\langle C \rangle$

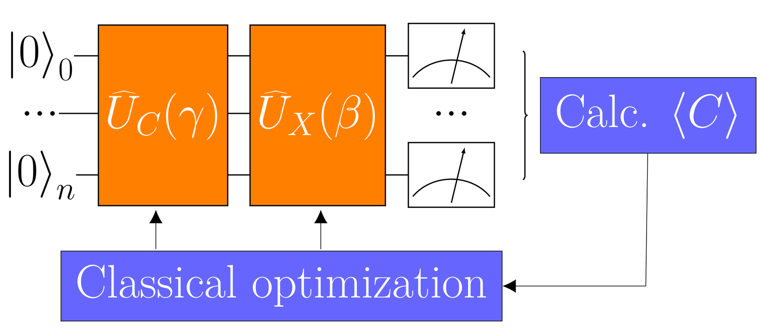

Quantum Annealing

- Determines minimum energy configuration
- Efficiently explores energy landscape using quantum phenomena (e.g., quantum tunneling)
- Systems with ≈ 5 000 qubits provided by D-Wave

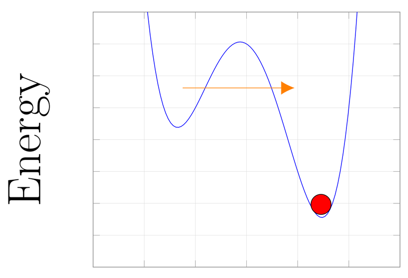

Challenges and Limitations

General Challenges

- Limited qubit numbers
- Problem encodings (e.g., QUBO) required
- Reproducing experiments [2]

Limited Qubit Connectivity

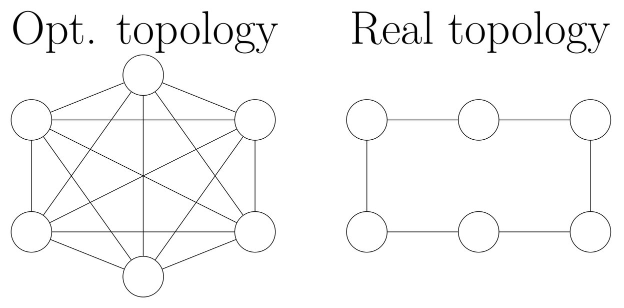

Limited Coherence Time

- Decoherence: Gradual decay of quantum states
- Deeper circuits increase execution time and chance of decoherence errors

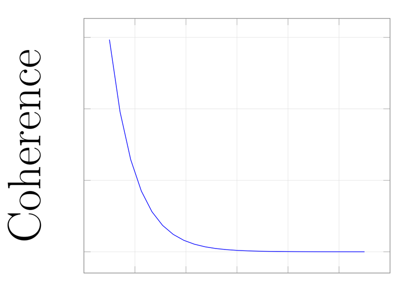

- The circuit depth is a crucial metric for quantum computing feasibility
- Minimize required qubits and circuit depth:
  - keep the encoding lightweight
  - reduce the number of qubit interactions

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 way: Choose locally cheapest plan
- Better: Select plans with common subexpressions

QUBO Reformulation

- Energy formula: $\omega_L E_L + \omega_M E_M + E_C + E_S$ [3]
- QUBO terms for incentivizing valid and optimal solutions:
  - $E_L = -\sum_{p \in P} X_p$
  - $E_M = \sum_{q \in Q} \sum_{\{p1,p2\} \subseteq P_q} X_{p1} X_{p2}$
  - $E_C = \sum_{p \in P} c_p X_p$
  - $E_S = -\sum_{\{p1,p2\} \subseteq P} s_{p1,\,p2} X_{p1} X_{p2}$

Solving MQO on IBM-Q QPUs

New!

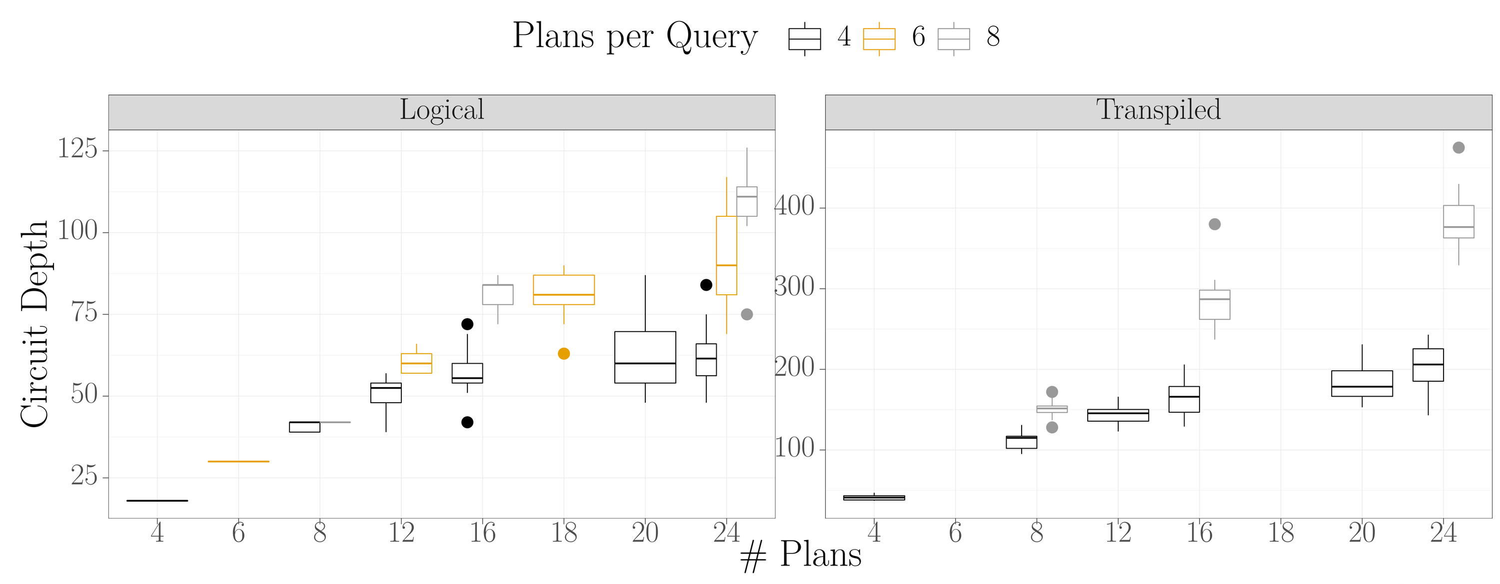

- Solving multi query optimization (MQO) with QAOA on gate-based QPUs [4]
- Significant depth increase when transpiling circuits

Join Ordering Reformulation

New!

Problem Classification

- General query graphs
- Left-deep join trees
- Support of cross products
- Min. intermediate cardinalities

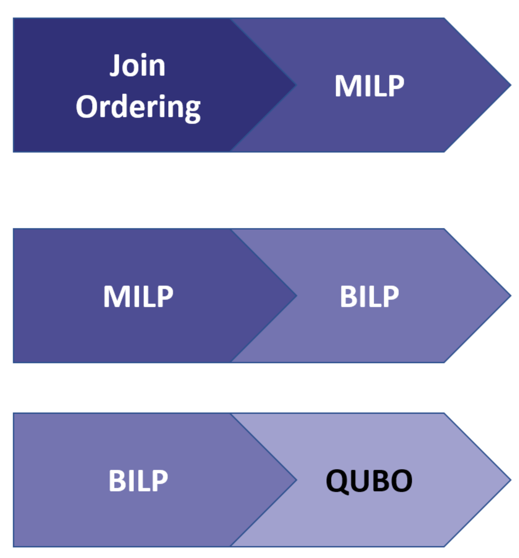

- Mixed integer linear programming (MILP) reformulation [5]
- Approximate log. cardinalities: $\min \sum_{r=0}^{R-1} \sum_{j=1}^{J-1} cto_{rj} \theta_r$
- Approx./ validity constraints: $c_j - cto_{rj} \cdot \infty_{rj} \leq \log(\theta_r)$, ...
- Equality conversion: $c_j - cto_{rj} \cdot \infty_{rj} + s_{rj} = \log(\theta_r)$
- Variable discretization: $s_{rj} \approx \omega \sum_{i=1}^{n} 2^{i-1} b_i$
- Transform the binary ILP (BILP) problem to QUBO [6]
- 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!

Findings

- IBM-Q QPUs so far only allow small scale queries due to qubit limitations
- D-Wave systems support queries joining up to 15 relations
- Large impact of an increased precision:
  - Higher qubit consumption
  - IBM-Q: Drastic depth increase
  - D-Wave: Significant reduction of problem sizes

IBM-Q Results

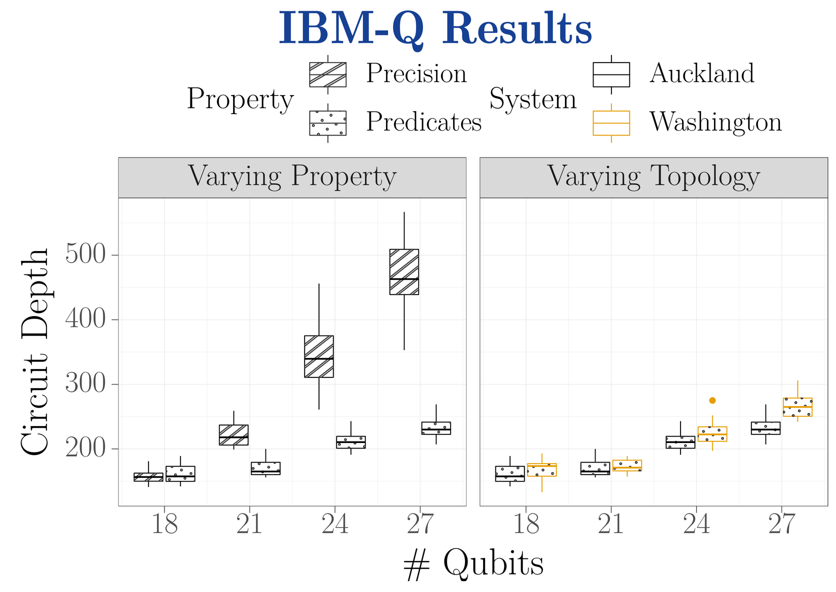

D-Wave Results

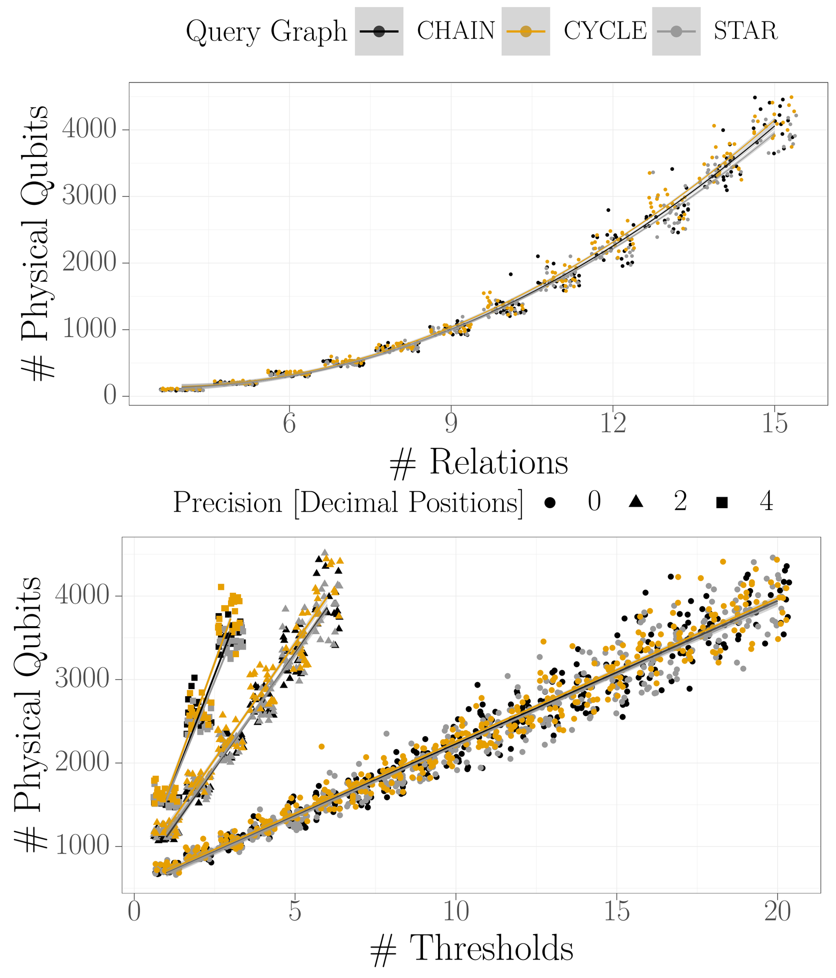

Path to DB-QPU Integration

- Non-traditional problem implementations (e.g., QUBO) required
- Subtle issues that are negligible for classical CPUs may have a large impact on quantum computing
- Qubit connectivity identified as a large bottleneck for join ordering
- Not only qubit limits, but all bottlenecks need to be addressed by 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.

[4] 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.

External Publications

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

[5] 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.

[6] 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.

I thank my supervisors Wolfgang Mauerer (Technical University of Applied Sciences Regensburg) and Stefanie Scherzinger (University of Passau) for their support.

Federal Ministry of Education and Research