How Can Quantum Computing and Databases Meet "in Practice"? From Algorithms to Systems

Hanwen Liu
hanwen_liu@usc.edu
University of Southern California
Los Angeles, California, USA

Yiran Li
one.li@utoronto.ca
University of Toronto
Toronto, Ontario, Canada

Ibrahim Sabek
sabek@usc.edu
University of Southern California
Los Angeles, California, USA

Xuanhe Zhou
zhouxuanhe@sjtu.edu.cn
Shanghai Jiao Tong University
Shanghai, China

Hans-Arno Jacobsen
jacobsen@eecg.toronto.edu
University of Toronto
Toronto, Ontario, Canada

ABSTRACT

Quantum computing has recently moved from a theoretical concept to a transformative and increasingly accessible computational paradigm. These developments sparked growing interest in the database community along two complementary directions: applying quantum techniques to improve database tasks, and leveraging DBMS principles to support quantum workloads. This three-hour tutorial provides a bidirectional, practice-oriented treatment of this intersection, covering quantum foundations, optimization algorithms, a systematic survey of Quantum-for-DB and DB-for-Quantum research, and end-to-end system integration with a real DBMS. The tutorial includes hands-on demonstrations using quantum simulators and real quantum hardware. We summarize open challenges and research opportunities for practical real-world deployment.

PVLDB Reference Format:
Hanwen Liu, Yiran Li, Ibrahim Sabek, Xuanhe Zhou, and Hans-Arno Jacobsen. How Can Quantum Computing and Databases Meet "in Practice"? From Algorithms to Systems. PVLDB, 19(12): 4906 - 4911, 2026.
doi:10.14778/3827998.3828148

PVLDB Artifact Availability:
The tutorial slides and demonstration code will be made publicly available at https://github.com/ihanwen99/vldb26-tutorial.

1 INTRODUCTION

Driven by rapid advances in hardware [42, 52, 53], quantum computing has emerged as a transformative and promising field. It spans four paradigms: gate-based quantum computing [67], annealing-based quantum computing [5], hybrid quantum-classical methods [41], and quantum-inspired methods [2]. These advances motivated a growing body of database research in two complementary directions. On the one hand, researchers have applied these quantum and quantum-inspired methods to a range of database tasks, including join order optimization [44, 46, 55–57, 64], multiple query optimization [10, 58, 61], index selection [22, 25, 62], transaction scheduling [3, 20, 43], and schema matching [13, 17]. On the other hand, researchers have also begun to explore how database systems and DBMS principles can support quantum workloads [16, 31, 59, 60, 63]. Taken together, these developments suggest that the intersection of quantum computing and databases is no longer speculative or far-fetched. However, the existing landscape remains fragmented, making it timely to ask not only what quantum computing can do for databases, or what databases can do for quantum workloads, but also how the two can meet *"in practice"*.

In this tutorial, we present a bidirectional view of the interplay between quantum computing and database systems: (i) how quantum computing can accelerate and extend database systems, and (ii) how database principles can enable efficient and scalable quantum workloads. Our focus goes beyond surveying algorithms to emphasize practical integration and co-design, highlighting key challenges and research opportunities to make quantum computing viable in real data systems. Prior tutorials at SIGMOD 2023 [65], ICDE 2024 [23], and SIGMOD 2026 [27] explored one-directional applications of quantum computing for database problems. A QCE 2024 tutorial [45] mostly focuses on quantum annealing. Our tutorial differs by systematically categorizing work across all four quantum computing paradigms, including quantum-inspired techniques on classical hardware; covering the DB-for-Quantum perspective as essential infrastructure for quantum workloads; and emphasizing practical systems integration with a real DBMS.

Target Audience. This tutorial targets researchers, developers, and practitioners interested in quantum computing and data management systems. No prior quantum background is required. For graduate students, the tutorial provides a structured overview of existing work and open problems, helping them identify promising research topics. Overall, the tutorial serves as an invitation to the database community to engage with quantum computing.

Relevance To VLDB. Research on quantum computing for data management has become increasingly active in the database community, particularly within VLDB. With the growing number of proposed approaches and systems, it is timely to offer a tutorial that surveys the current state of the art and outlines promising future research directions. Moreover, recent work has demonstrated practical quantum-augmented database optimization on real hardware and real DBMS workloads, making a practice-oriented, hands-on tutorial particularly timely. Many of the efforts covered in this

Proceedings of the VLDB Endowment, Vol. 19, No. 12 ISSN 2150-8097.
doi:10.14778/3827998.3828148

tutorial have been recently published in major database venues, including VLDB, SIGMOD, and ICDE.

2 TUTORIAL OUTLINE

This three-hour tutorial is organized into the following six parts. We provide a summary of representative works in Table 1. Part 2 and Part 5 contain hands-on demonstration sessions.

Part 1.	Quantum Background & Motivation	20 min
Part 2.	Quantum Computing Paradigms & Algorithms	30 min
Part 3.	Quantum Computing for Database Systems (DB)	40 min
Part 4.	Data Management for Quantum Computing (QC)	30 min
Part 5.	QC & DB Meet "in Practice"	40 min
Part 6.	Open Challenges & Research Opportunities	20 min

2.1 Quantum Background and Motivation

We begin with the computational foundations of quantum computing: qubits and quantum states, quantum gates and circuits, and the key phenomena of superposition, entanglement, and decoherence, presented at an operational level without requiring prior knowledge of quantum mechanics. We then introduce the Quadratic Unconstrained Binary Optimization (QUBO) [24] framework and the Ising model [40] as the formulation bridges between classical optimization problems and quantum algorithms. These formulations provide a unified interface through which database optimization tasks are encoded for quantum solvers.

2.2 Computing Paradigms and Algorithms

This part introduces four quantum computing paradigms and their core algorithms for solving the problems formulated above.

Quantum Annealing (QA) [5] targets discrete optimization by evolving a physical system toward low-energy solutions, often leveraging quantum tunneling to escape poor local minima [54]. QA is based on the principles of the adiabatic quantum computing model [1, 12], but specializes them to combinatorial optimization by encoding problems as Ising models or QUBOs.

Gate-Based Quantum Computing [67] applies sequences of unitary quantum gates to qubits, providing a universal model for quantum algorithms. Grover's search [21], with adaptive variants [18], is an important quantum method for solving unstructured search problems with a quadratic speedup over classical algorithms.

Hybrid Quantum-Classical Approaches [41] combine quantum methods with classical orchestration, optimization, and post-processing, and are particularly important in the NISQ era where purely quantum execution remains resource-constrained [9]. The Quantum Approximate Optimization Algorithm (QAOA) [11] is a prominent example, using parameterized quantum circuits whose parameters are iteratively updated by a classical optimizer.

Quantum-Inspired Approaches [2] do not rely on quantum hardware, but instead use classical algorithms or specialized accelerators (e.g., Fujitsu Digital Annealer [14]) motivated by quantum principles to address hard combinatorial optimization problems [48].

2.3 Quantum Computing for Database Systems

This part of the tutorial examines how quantum computing can help database systems in practice. We first review optimization-centric work, where quantum or quantum-inspired methods are applied to database optimization problems. We then broaden the view to more general use cases beyond optimization.

Optimization-Centric Opportunities. Most related work in this line concentrates on using quantum computing to solve database optimization problems. This part of the tutorial organizes the literature by optimization task and, within each task, highlights the corresponding quantum computing paradigms.

Join Order Optimization. This is one of the most extensively explored optimization problems in the database community. Most works formulate it as a QUBO problem [40] and solve it using a range of quantum (and quantum-inspired) paradigms, including quantum annealing [44, 56], gate-based models [44], quantum-inspired digital annealers [14, 57], and hybrid quantum-classical solvers [8, 38, 47, 55]. There is also a growing line of work on quantum machine learning for join order optimization [64]. These studies suggest that join order optimization is a suitable application and testbed for various quantum methods. A key practical gap is that most of these approaches do not evaluate the resulting join orders in an actual database system. Note that improved objective values in the quantum formulation (e.g., lower QUBO energy) do not always translate into lower query latency in practice.

Multiple Query Optimization (MQO). MQO aims to reduce redundant work during the execution of a batch of queries, yielding substantial performance benefits in database systems. Early work [61] introduced a simplified formulation and executed it on a quantum annealer, providing one of the first experimental studies in this space. More recent efforts propose hybrid quantum-classical approaches to improve efficiency and scalability [10]. Most recently, incremental quantum-inspired annealing methods have been developed to target larger-scale MQO settings [58]. A key practical gap is that existing MQO formulations are often overly simplified, so solutions produced by these quantum approaches do not directly translate into deployable solutions in real database systems.

Index Selection. Index selection is another NP-hard [7] database optimization problem that has been explored with quantum annealing [62], gate-based quantum computing [25], and quantum machine learning [22]. A key practical gap is that these approaches are typically not integrated into, and evaluated within, actual databases.

Transaction Scheduling. Transaction scheduling focuses on determining an execution order for concurrent transactions. Prior work has explored this problem using quantum annealing [3], Grover's search algorithm [20], and a hybrid quantum-classical approach that iteratively uses a locking mechanism [43]. Although they reduce formulation complexity as a trade-off to improve solvability on existing devices, scalability remains less discussed.

Schema Matching. In its general form, this problem is NP-hard. It has been addressed with quantum annealing [13] and QAOA [17].

Graph Analytics. NP-hard graph problems often arise in optimizing data placement and scheduling. Recent work has explored Grover's search [29] and quantum optimization methods [30].

Beyond Optimization: Retrieval and Execution. Theoretical foundations, such as quantum algorithms for element distinctness [4], suggest that gate-based quantum computers can perform database-like operations (e.g., multi-element retrieval and deduplication) faster than classical approaches. Building on theoretical improvements in computational complexity, recent studies have

Table 1: Comparison of Representative Works among Quantum Computing and Database Systems.

Types	Representative Works	Problems	Computing Paradigms				Key Practical Gaps
			Gate-based	Annealing-based	Hybrid Quantum-Classical	Quantum-Inspired	
QC4DB	[44, 46, 55–57, 64]	Join Order Optimization	✓	✓	✓	✓	Limited DBMS Validation
	[10, 58, 61]	Multiple Query Optimization		✓	✓	✓	Simplified Formulation
	[22, 25, 62]	Index Selection	✓	✓	✓		No DBMS Integration
	[3, 20, 43]	Transaction Scheduling	✓	✓	✓		Scalability Unexplored
	[13, 17]	Schema Matching		✓	✓		No Pipeline Integration
	[4, 26, 32]	Retrieval and Execution	✓				Overhead
	[66]	Approximate Query Processing			✓		Scalability
	[15]	Privacy	✓				Hardware Feasibility
	[29, 30]	Graph analytics	✓	✓	✓		No Hardware Validation
DB4QC	[19, 28]	Storage and Encoding	✓				Measurement Collapse
	[16, 31, 59, 63]	Simulation and Transpilation	✓				Performance Gap
Bidirectional	[33–35, 37, 39]	Actual Integration			✓		Generalizability
	[6, 26, 51]	Quantum-Native DB Arch.	✓		✓		Efficient SQL-for-Gate

investigated the viability of quantum-based SQL execution by mapping relational operators (selection σ, join $\bowtie$, union $\cup$, intersection $\cap$, difference $-$, and sort) to Grover oracles within a gate-based circuit model [26]. Quantum Partitioned Databases (QPD) are designed specifically to enhance database performance by reducing reliance on classical B-trees, using modified Grover algorithms for multi-element retrieval on gate-based quantum processors [32].

Quantum-Enhanced Approximate Query Processing. Approximate query processing (AQP) answers aggregate queries over large datasets with bounded error, and a major challenge is estimating aggregates for rare, long-tail groups where random sampling provides too few representatives. A hybrid quantum-classical approach that combines classical and gate-based quantum sampling addresses this issue [66]: classical sampling first identifies underrepresented groups, and gate-based quantum amplitude amplification selectively increases their sampling probability before measurement. The amplified samples are then corrected via classical post-processing. However, the scalability challenge remains significant: the number of amplification rounds grows with the inverse of the group's classical probability, and the QRAM encoding required to load the dataset into quantum states introduces substantial overhead.

Quantum Privacy for Databases. Beyond speedups, quantum mechanics offers primitives with no direct classical analogue. For example, the no-cloning principle can be leveraged for database privacy. Recent work implements Symmetric Private Information Retrieval (SPIR) on NISQ gate-based hardware using Quantum Random Access Codes (QRAC) and mutually unbiased bases [15]. In this architecture, the classical DBMS stores compact state descriptions and, upon a query, prepares fresh quantum states for the client to measure. The protocol hides the user's query from the server while restricting the client to learning only the requested tuple; attempts to copy or extract additional information are inherently limited and become detectable through measurement disturbance. The critical challenge lies in the practical feasibility on current NISQ devices. The copy budget k (the number of fresh quantum states generated per query) directly governs the trade-off between retrieval correctness and privacy strength.

2.4 Data Management for Quantum Computing

This part of the tutorial offers the reverse perspective: Instead of asking how quantum computing can help database systems, we ask how database systems can support quantum workloads.

Quantum Data Storage and Encoding. The fundamental differences between quantum and classical computing systems are in data representations. In classical databases, data is stored deterministically. In quantum systems, data is typically encoded into qubits and manipulated via quantum gates, often exploiting superposition to represent multiple states simultaneously. Recent efforts have explored quantum storage designs for relational tables. For example, QCOS (column-oriented) and QROS (row-oriented) encode tables into superposition states within quantum circuits [28]. Using Hadamard and multi-controlled Toffoli (MCT) gates, these designs aim to achieve logarithmic qubit scaling with respect to data volume. More broadly, different problem encoding patterns induce different trade-offs in efficiency and scalability [19]. However, a significant challenge remains: measurement collapse. Once a quantum state is measured to retrieve data, the superposition collapses, and the stored information is destroyed. As a result, the state must be re-prepared for subsequent queries, introducing loading overhead.

DB for Quantum Simulation and Transpilation. Since large-scale, fault-tolerant quantum computers are not yet widely available, simulating quantum computation on classical hardware remains essential for algorithm development and quantum hardware validation. A core challenge is that simulation memory consumption grows exponentially with the number of qubits. To address this, researchers have mapped quantum states and gate operations to SQL queries [59] and executed them natively inside relational database engines [31]. More recently, work has explored evaluating conjunctive queries to simulate quantum circuits [16] and leveraging graph-database techniques for circuit transpilation [63]. However, a key challenge is the performance gap between SQL-based tensor contractions and specialized high-performance computing (HPC) tensor-network libraries. Optimizing complex, dense quantum circuits (e.g., the Quantum Fourier Transform) within a relational engine remains a formidable obstacle.

2.5 QC and DB Meet "in Practice"

This part of the tutorial focuses on how quantum computing and database systems can meet in practical settings.

Actual Integration. We highlight recent efforts to integrate a quantum computing component into actual database systems. The representative line of work is *Q^2O* [33–35, 39], the first quantum-augmented query optimizer integrated with PostgreSQL (an *actual database system*) and executed on *hybrid quantum-classical solvers*, demonstrating that quantum computing can improve the performance of *real workloads*. To extend beyond prior simplified problem settings restricted to left-deep structures, they designed an efficient quantum algorithm for complex bushy join trees [33]. Moreover, the end-to-end latency improvement provides concrete evidence [35, 37] that the advantages of quantum computing can be realized in broader real-time applications.

Quantum-Native Database Architecture. Beyond integration, an emerging direction is co-design, which rethinks database system architecture in the presence of quantum computation. Kesarwani et al. [26] brought this direction into systems focus by examining relational query execution on quantum platforms. Subsequent work argues that a quantum database must support a full lifecycle of operations, including preparation, extension, writing, and reading, while respecting quantum-mechanical constraints such as the no-cloning theorem and the inability to arbitrarily delete unknown quantum states [51]. Building on this view, recent work has proposed "quantum–classical" settings in which quantum indexing is used alongside classical data; for example, an index register requires only $\lceil \log_2(k) \rceil$ qubits to represent k index elements in superposition [51]. A central challenge in these systems is bridging declarative database operations (e.g., SQL) with low-level quantum gate execution. Pushing this direction further, *Qute* [6] explores a quantum-native database architecture that employs hybrid optimizers to select between classical and quantum execution plans while operating under the strict qubit budgets of current NISQ devices.

2.6 Hands-on Demonstrations

Algorithmic Formulation on Simulators (Part 2). We include a QUBO-based optimization example, illustrating how a combinatorial optimization problem can be formulated in a way amenable to quantum computational workflows, and how to retrieve results from a simulation of QAOA-based optimization. We also demonstrate Grover's search, with emphasis on the general search formulation and the role of the oracle. These demonstrations are delivered via Jupyter notebooks on quantum simulators, so that participants can run them on local machines without requiring any external hardware access.

Practical Quantum Computing for Databases (Part 5). We prepare two demonstration scenarios, running on real quantum annealing hardware. The first is *Q^2O* [39], a quantum-augmented query optimizer integrated with PostgreSQL. The second is *QFusion* [36], a boundary-aware fusion framework that scales QUBO optimization to large database optimization problems.

2.7 Research Challenges and Opportunities

The central challenge now is to systematically connect quantum applications, algorithms, and systems for practical, controllable, scalable, and manageable real-world deployment. From this perspective, we summarize four open questions: (Q1) What limits practical representations and execution? (Q2) How can controllable algorithms balance solution quality and efficiency? (Q3) Why is scalability critical for real-world quantum applications? (Q4) Why do quantum workflows need database-style storage management?

3 PRESENTERS

Hanwen Liu is a PhD candidate in the Department of Computer Science at the University of Southern California. He received his M.Sc. degree from the Technical University of Munich in 2024. His research interests lie at the intersection of scalable quantum computing and efficient data management systems. During his PhD, he collaborated with the AWS Learned Systems Group. Hanwen has published multiple papers at VLDB, SIGMOD, and ICDE.

Yiran Li is a postdoctoral fellow in the Department of Electrical and Computer Engineering at the University of Toronto, where he works with Hans-Arno Jacobsen. He received his Ph.D. from The Hong Kong Polytechnic University, where his research focused on clustering and embedding algorithms for large-scale graph data. His research interests include graph data processing and analytics, and the intersection of data management with quantum computing.

Ibrahim Sabek is an Assistant Professor of Computer Science at the University of Southern California. His current research interests focus on building the next generation of data management, processing, and analysis systems using quantum computing and machine learning. His research work has been recognized by the 2026 NSF CAREER Award, the 2025 Google ML and Systems Junior Faculty Award, the 2024 and 2025 Google Data Analytics and Insights (DANI) Awards, the NSF Computing Innovation Fellowship (2020-2023), among others. Ibrahim is a workshop co-chair for the quantum computing workshops in the database and spatial computing communities, namely, Q-Data@SIGMOD [49] and Q-Spatial@SIGSPATIAL [50].

Xuanhe Zhou is an Assistant Professor in the Department of Computer Science, Shanghai Jiao Tong University. His research interests lie in data systems. More recently, he has worked on quantum-native database systems. He has received the SIGMOD 2025 Jim Gray Dissertation Honorable Mention and VLDB 2023 Best Industry Paper Runner-up.

Hans-Arno Jacobsen holds the Jeffrey Skoll Chair in Computer Networking and Innovation and is a Tier 1 Canada Research Chair in Data-Intensive Systems at the University of Toronto, where he is Professor of Electrical and Computer Engineering and Computer Science. His research spans distributed systems, data management, and AI-driven systems, with a focus on large-scale, data-intensive infrastructures for cloud, telecom, and emerging computing paradigms. He has made foundational contributions to event processing and scalable distributed systems. Arno is also advancing research in quantum computing, including quantum machine learning, distributed quantum systems, and applications in computational chemistry. He serves as Scientific Director of the pan-Canadian Quantum Software Consortium (`qscc.ca`). More information on his current and past research can also be found at `msrg.org`. He is a Fellow of the IEEE.

REFERENCES

[1] Tameem Albash and Daniel A. Lidar. 2018. Adiabatic Quantum Computation. *Reviews of Modern Physics* 90, 1 (Jan. 2018), 015002.

[2] Juan Miguel Arrazola, Alain Delgado, Bhaskar Roy Bardhan, and Seth Lloyd. 2020. Quantum-inspired algorithms in practice. *Quantum* 4 (2020), 307. https://doi.org/10.22331/q-2020-08-13-307

[3] Tim Bittner and Sven Groppe. 2020. Avoiding blocking by scheduling transactions using quantum annealing. In *Proceedings of the 24th Symposium on International Database Engineering & Applications* (Seoul, Republic of Korea) *(IDEAS '20)*. Association for Computing Machinery, New York, NY, USA, Article 21, 10 pages. https://doi.org/10.1145/3410566.3410593

[4] Harry Buhrman, Christoph Dürr, Mark Heiligman, Peter Høyer, Frédéric Magniez, Miklos Santha, and Ronald de Wolf. 2005. Quantum Algorithms for Element Distinctness. *SIAM J. Comput.* 34, 6 (2005), 1324–1330.

[5] Umut Çalikyilmaz, Sven Groppe, Jinghua Groppe, Tobias Winker, Stefan Prestel, Farida Shagieva, Daanish Arya, Florian Preis, and Le Gruenwald. 2023. Opportunities for Quantum Acceleration of Databases: Optimization of Queries and Transaction Schedules. *Proc. VLDB Endow.* 16, 9 (May 2023), 2344–2353. https://doi.org/10.14778/3598581.3598603

[6] Muzhi Chen, Xuanhe Zhou, Wei Zhou, Bangrui Xu, Surui Tang, Guoliang Li, Bingsheng He, Yeye He, Yitong Song, and Fan Wu. 2026. Qute: Towards Quantum-Native Database. *CoRR* abs/2602.14699 (2026).

[7] Douglas Comer. 1978. The Difficulty of Optimum Index Selection. *ACM Trans. Database Syst.* 3, 4 (1978), 440–445. https://doi.org/10.1145/320289.320296

[8] D-Wave Systems Inc. 2021. Hybrid Solver for Constrained Quadratic Models. https://www.dwavesys.com/media/rldh2ghw/14-1055a-a_hybrid_solver_for_constrained_quadratic_models.pdf. White paper. Accessed: 2026-08-07.

[9] Suguru Endo, Zhenyu Cai, Simon C. Benjamin, and Xiao Yuan. 2021. Hybrid Quantum-Classical Algorithms and Quantum Error Mitigation. *Journal of the Physical Society of Japan* 90, 3 (2021), 032001. https://doi.org/10.7566/JPSJ.90.032001

[10] Tobias Fankhauser, Marc E. Solèr, Rudolf M. Füchslin, and Kurt Stockinger. 2021. Multiple Query Optimization using a Hybrid Approach of Classical and Quantum Computing. arXiv:2107.10508 [cs.DB]

[11] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. 2014. A Quantum Approximate Optimization Algorithm. arXiv:1411.4028 [quant-ph]

[12] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. 2000. Quantum Computation by Adiabatic Evolution. arXiv:quant-ph/0001106

[13] Kristin Fritsch and Stefanie Scherzinger. 2023. Solving Hard Variants of Database Schema Matching on Quantum Computers. *Proc. VLDB Endow.* 16, 12 (Aug. 2023), 3990–3993. https://doi.org/10.14778/3611540.3611603

[14] Fujitsu Limited. 2023. Fujitsu Digital Annealer - solving large-scale combinatorial optimization problems instantly. https://www.fujitsu.com/global/services/business-services/digital-annealer/what-is-digital-annealer/.

[15] Giancarlo Gatti, Floris Geerts, and Rihan Hai. 2025. Private Quantum Database. *CoRR* abs/2508.19055 (2025).

[16] Floris Geerts and Rihan Hai. 2025. QC meet CQ: Quantum Conjunctive Queries. In *Proceedings of the 2nd Workshop on Quantum Computing and Quantum-Inspired Technology for Data-Intensive Systems and Applications (Q-Data '25)*. Association for Computing Machinery, New York, NY, USA, 25. https://doi.org/10.1145/3736393.3736696

[17] Luisa Gerlach, Tobias Köppl, Stefanie Scherzinger, Nicole Schweikardt, and René Zander. 2025. A Quantum-Leap into Schema Matching: Beyond 1-to-1 Matchings. *Proc. ACM Manag. Data* 3, 2, Article 89 (June 2025), 24 pages. https://doi.org/10.1145/3725226

[18] Austin Gilliam, Stefan Woerner, and Constantin Gonciulea. 2021. Grover Adaptive Search for Constrained Polynomial Binary Optimization. *Quantum* 5 (April 2021), 428. https://doi.org/10.22331/q-2021-04-08-428 arXiv:1912.04088 [quant-ph]

[19] Martin Gogeissl, Hila Safi, and Wolfgang Mauerer. 2024. Quantum Data Encoding Patterns and their Consequences. In *Proceedings of the 1st Workshop on Quantum Computing and Quantum-Inspired Technology for Data-Intensive Systems and Applications* (Santiago, Chile) *(Q-Data '24)*. Association for Computing Machinery, New York, NY, USA, 27–37. https://doi.org/10.1145/3665225.3665446

[20] Sven Groppe and Jinghua Groppe. 2021. Optimizing Transaction Schedules on Universal Quantum Computers via Code Generation for Grover's Search Algorithm. In *Proceedings of the 25th International Database Engineering & Applications Symposium* (Montreal, QC, Canada) *(IDEAS '21)*. Association for Computing Machinery, New York, NY, USA, 149–156. https://doi.org/10.1145/3472163.3472164

[21] Lov K. Grover. 1996. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing* (Philadelphia, Pennsylvania, USA) *(STOC '96)*. Association for Computing Machinery, New York, NY, USA, 212–219. https://doi.org/10.1145/237814.237866

[22] Le Gruenwald, Tobias Winker, Umut Çalikyilmaz, Jinghua Groppe, and Sven Groppe. 2023. Index Tuning with Machine Learning on Quantum Computers for Large-Scale Database Applications. In *Joint Proceedings of Workshops at the 49th International Conference on Very Large Data Bases (VLDB 2023), Vancouver, Canada, August 28 - September 1, 2023 (CEUR Workshop Proceedings)*, Vol. 3462. CEUR-WS.org. https://ceur-ws.org/Vol-3462/QDSM5.pdf

[23] Rihan Hai, Shih-Han Hung, and Sebastian Feld. 2024. Quantum Data Management: From Theory to Opportunities. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. IEEE, 5376–5381. https://doi.org/10.1109/ICDE60146.2024.00410

[24] Joseph T. Iosue. 2019. qubovert Documentation. https://qubovert.readthedocs.io/en/latest/. Accessed: 2026-08-07.

[25] Manish Kesarwani and Jayant R. Haritsa. 2024. Index Advisors on Quantum Platforms. *Proc. VLDB Endow.* 17, 11 (July 2024), 3615–3628. https://doi.org/10.14778/3681954.3682025

[26] Manish Kesarwani and Jayant R. Haritsa. 2024. Is Quantum-Based SQL Query Execution Viable?. In *Proceedings of Workshops at the 50th International Conference on Very Large Data Bases, VLDB 2024, Guangzhou, China, August 26-30, 2024*. VLDB.org.

[27] Manish Kesarwani and Jayant R. Haritsa. 2026. Relational Database Engines on Quantum Platforms: Concepts, Algorithms, and Implementations. In *Companion of the International Conference on Management of Data* (Bengaluru, India) *(SIGMOD Companion '26)*. Association for Computing Machinery, New York, NY, USA, 556–562. https://doi.org/10.1145/3788853.3801876

[28] Tuodu Li, Gongsheng Yuan, Chang Yao, Meng Shi, Ziyue Wang, Ling Qian, and Jiaheng Lu. 2024. Quantum Storage Design for Tables in RDBMS. In *Proceedings of Workshops at the 50th International Conference on Very Large Data Bases, VLDB 2024, Guangzhou, China, August 26-30, 2024*. VLDB.org. https://vldb.org/workshops/2024/proceedings/QDSM/QDSM.4.pdf

[29] Xiaofan Li, Gao Cong, and Rui Zhou. 2024. Quantum Algorithms for the Maximum K-Plex Problem. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. 2435–2448.

[30] Yiran Li, Y. Batuhan Yilmaz, Michael Silver, Zachary Vernec, and Hans-Arno Jacobsen. 2026. Quantum Hypergraph Partitioning. arXiv:2605.02635 [cs]

[31] Tim Littau and Rihan Hai. 2025. Qymera: Simulating Quantum Circuits using RDBMS. In *Companion of the 2025 International Conference on Management of Data* (Berlin, Germany). ACM, 179–182. https://doi.org/10.1145/3722212.3725126

[32] Tim Littau, Ziyu Li, and Rihan Hai. 2024. Towards Quantum Data Structures for Enhanced Database Performance. In *VLDB 2024 Workshop: The Second International Workshop on Quantum Data Science and Management (QDSM'24)*.

[33] Hanwen Liu, Abhishek Kumar, Federico Spedalieri, and Ibrahim Sabek. 2025. Hybrid Quantum-Classical Optimization for Bushy Join Trees. In *Proceedings of the 2nd Workshop on Quantum Computing and Quantum-Inspired Technology for Data-Intensive Systems and Applications (Q-Data '25)*. Association for Computing Machinery, New York, NY, USA, 20–24. https://doi.org/10.1145/3736393.3736695

[34] Hanwen Liu, Abhishek Kumar, Federico Spedalieri, and Ibrahim Sabek. 2026. QDBO: A Real-time Quantum-augmented Database System Optimizer. *Proc. VLDB Endow.* 19, 11 (2026), 3606–3620. https://doi.org/10.14778/3836663.3836712

[35] Hanwen Liu and Ibrahim Sabek. 2026. Is Quantum Computing Ready for Real-Time Database Optimization?. In *2026 IEEE 42nd International Conference on Data Engineering (ICDE)*. https://arxiv.org/abs/2601.12123

[36] Hanwen Liu and Ibrahim Sabek. 2026. QFusion: A Demonstration of Boundary-Aware Fusion Planning and Execution for Large-Scale QUBO Optimization. *Proc. VLDB Endow.* 19, 12 (2026), 4826–4829. https://doi.org/10.14778/3827998.3828132

[37] Hanwen Liu and Ibrahim Sabek. 2026. Towards a Hybrid Quantum-Classical Computing Framework for Database Optimization Problems in Real Time Setup. In *ICDE 2026 Data Engineering Future Technologies*. https://arxiv.org/abs/2602.14263

[38] Hanwen Liu, Pranshi Saxena, Federico Spedalieri, and Ibrahim Sabek. 2025. Optimizing Join Orders via Constrained Quadratic Models (Abstract). In *Proc. 1st Workshop Quantum Data and Machine Learning* (Hong Kong, China) *(QDML '25)*.

[39] Hanwen Liu, Federico Spedalieri, and Ibrahim Sabek. 2025. A Demonstration of Q2O: Quantum-augmented Query Optimizer. *Proc. VLDB Endow.* 18, 12 (2025), 5439–5443. https://doi.org/10.14778/3750601.3750691

[40] Andrew Lucas. 2014. Ising formulations of many NP problems. *Frontiers in Physics* 2 (2014), 5. https://doi.org/10.3389/fphy.2014.00005

[41] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. 2016. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics* 18, 2 (2016), 023023. https://doi.org/10.1088/1367-2630/18/2/023023

[42] Dan Milmo. 2025. Google hails breakthrough as quantum computer surpasses ability of supercomputers. https://www.theguardian.com/technology/2025/oct/22/google-hails-breakthrough-as-quantum-computer-surpasses-ability-of-supercomputers. The Guardian. Accessed: 2026-08-07.

[43] Nitin Nayak, Alexandru Prisacaru, Umut Çalıkyılmaz, Jinghua Groppe, and Sven Groppe. 2025. Quantum-Enhanced Transaction Scheduling with Reduced Complexity via Solving QUBO Iteratively using a Locking Mechanism. In *Proceedings of the 2nd Workshop on Quantum Computing and Quantum-Inspired Technology for Data-Intensive Systems and Applications (Q-Data '25)*. Association for Computing Machinery, New York, NY, USA, 26–35. https://doi.org/10.1145/3736393.3736701

[44] Nitin Nayak, Jan Rehfeld, Tobias Winker, Benjamin Warnke, Umut Çalıkyılmaz, and Sven Groppe. 2023. Constructing Optimal Bushy Join Trees by Solving

QUBO Problems on Quantum Hardware and Simulators. In *Proceedings of the International Workshop on Big Data in Emergent Distributed Environments (BiDEDE '23)*. Article 7. https://doi.org/10.1145/3579142.3594298
[45] Nitin Nayak, Manuel Schönberger, Valter Uotila, Zhengtong Yan, Sven Groppe, Jiaheng Lu, and Wolfgang Mauerer. 2024. QCE'24 Tutorial: Quantum Annealing – Emerging Exploration for Database Optimization. arXiv:2411.04638 [quant-ph]
[46] Nitin Nayak, Tobias Winker, Umut Çalıkyılmaz, Sven Groppe, and Jinghua Groppe. 2024. Quantum Join Ordering by Splitting the Search Space of QUBO Problems. *Datenbank-Spektrum* 24, 1 (2024), 21–32. https://doi.org/10.1007/s13222-024-00468-3
[47] Eneko Osaba and Pablo Miranda-Rodriguez. 2025. D-Wave's Nonlinear-Program Hybrid Solver: Description and Performance Analysis. *IEEE Access* 13 (2025), 4724–4736. https://doi.org/10.1109/ACCESS.2025.3525620
[48] Hiroki Oshiyama and Masayuki Ohzeki. 2022. Benchmark of quantum-inspired heuristic solvers for quadratic unconstrained binary optimization. *Sci. Rep.* 12, 1 (2022), 2146. https://doi.org/10.1038/s41598-022-06070-5 arXiv:2104.14096 [quant-ph]
[49] Q-Data Workshop 2026. Q-Data: Workshop on Quantum Computing and Quantum-Inspired Technology for Data-Intensive Systems and Applications. https://www.qdataworkshop.com/. Collocated with ACM SIGMOD.
[50] Q-Spatial Workshop 2026. Q-Spatial: International Workshop on Quantum Computing and Quantum-Inspired Technologies for Spatial Data Systems and Applications. http://qspatial.github.io/. Collocated with ACM SIGSPATIAL.
[51] Carla Rieger, Michele Grossi, Gian Giacomo Guerreschi, Sofia Vallecorsa, and Martin Werner. 2024. Operational Framework for a Quantum Database. *CoRR* abs/2405.14947 (2024).
[52] Rigetti Computing. 2022. Rigetti quantum processors. https://qcs.rigetti.com/qpus.
[53] James Sanders. 2019. TechRepublic: D-Wave Announces 5,000-Qubit Fifth Generation Quantum Annealer. https://www.dwavequantum.com/company/newsroom/media-coverage/techrepublic-d-wave-announces-5-000-qubit-fifth-generation-quantum-annealer/. Accessed: 2026-08-07.
[54] Giuseppe E. Santoro, Roman Martoňák, Erio Tosatti, and Roberto Car. 2002. Theory of Quantum Annealing of an Ising Spin Glass. *Science* 295, 5564 (2002), 2427–2430. https://doi.org/10.1126/science.1068774 arXiv:https://www.science.org/doi/pdf/10.1126/science.1068774
[55] Pranshi Saxena, Ibrahim Sabek, and Federico Spedalieri. 2024. Constrained Quadratic Model for Optimizing Join Orders. In *Proceedings of the 1st Workshop on Quantum Computing and Quantum-Inspired Technology for Data-Intensive Systems and Applications* (Santiago, Chile) *(Q-Data '24)*. 38–44. https://doi.org/10.1145/3665225.3665447
[56] Manuel Schönberger, Stefanie Scherzinger, and Wolfgang Mauerer. 2023. Ready to Leap (by Co-Design)? Join Order Optimisation on Quantum Hardware. *Proc. ACM Manag. Data* 1, 1, Article 92 (may 2023), 27 pages. https://doi.org/10.1145/3588946
[57] Manuel Schönberger, Immanuel Trummer, and Wolfgang Mauerer. 2023. Quantum-Inspired Digital Annealing for Join Ordering. *Proc. VLDB Endow.* 17, 3 (Nov. 2023), 511–524. https://doi.org/10.14778/3632093.3632112
[58] Manuel Schönberger, Immanuel Trummer, and Wolfgang Mauerer. 2025. Large-Scale Multiple Query Optimisation with Incremental Quantum(-Inspired) Annealing. *Proc. ACM Manag. Data* 3, 4 (2025), 253:1–253:25. https://doi.org/10.1145/3749171
[59] Immanuel Trummer. 2024. Towards Out-of-Core Simulators for Quantum Computing. In *Proceedings of the 1st Workshop on Quantum Computing and Quantum-Inspired Technology for Data-Intensive Systems and Applications* (Santiago, Chile) *(Q-Data '24)*. Association for Computing Machinery, New York, NY, USA, 1–1. https://doi.org/10.1145/3665225.3665441
[60] Immanuel Trummer. 2025. Cost-Based Query Optimization for Quantum Computation. In *Proceedings of the 2nd Workshop on Quantum Computing and Quantum-Inspired Technology for Data-Intensive Systems and Applications (Q-Data '25)*. Association for Computing Machinery, New York, NY, USA, 1–2. https://doi.org/10.1145/3736393.3736690
[61] Immanuel Trummer and Christoph Koch. 2016. Multiple Query Optimization on the D-Wave 2X Adiabatic Quantum Computer. *Proc. VLDB Endow.* 9, 9 (May 2016), 648–659. https://doi.org/10.14778/2947618.2947621
[62] Immanuel Trummer and Davide Venturelli. 2024. Leveraging Quantum Computing for Database Index Selection. In *Proceedings of the 1st Workshop on Quantum Computing and Quantum-Inspired Technology for Data-Intensive Systems and Applications* (Santiago, Chile) *(Q-Data '24)*. Association for Computing Machinery, New York, NY, USA, 14–26. https://doi.org/10.1145/3665225.3665445
[63] Valter Johan Edvard Uotila. 2025. Transpiling Quantum Circuits on Graph Databases. In *Proceedings of the 2nd Workshop on Quantum Computing and Quantum-Inspired Technology for Data-Intensive Systems and Applications (Q-Data '25)*. Association for Computing Machinery, New York, NY, USA, 18–19. https://doi.org/10.1145/3736393.3736694
[64] Tobias Winker, Umut Çalikyilmaz, Le Gruenwald, and Sven Groppe. 2023. Quantum Machine Learning for Join Order Optimization using Variational Quantum Circuits. In *Proceedings of the International Workshop on Big Data in Emergent Distributed Environments* (Seattle, WA, USA) *(BiDEDE '23)*. Association for Computing Machinery, New York, NY, USA, Article 5, 7 pages. https://doi.org/10.1145/3579142.3594299
[65] Tobias Winker, Sven Groppe, Valter Uotila, Zhengtong Yan, Jiaheng Lu, Maja Franz, and Wolfgang Mauerer. 2023. Quantum Machine Learning: Foundation, New Techniques, and Opportunities for Database Research. In *Companion of the 2023 International Conference on Management of Data* (Seattle, WA, USA) *(SIGMOD '23)*. Association for Computing Machinery, New York, NY, USA, 45–52. https://doi.org/10.1145/3555041.3589404
[66] Sai Wu, Meng Shi, Dongxiang Zhang, Junbo Zhao, Gongsheng Yuan, and Gang Chen. 2024. When Quantum Computing Meets Database: A Hybrid Sampling Framework for Approximate Query Processing. *IEEE Trans. Knowl. Data Eng.* 36, 12 (2024), 9532–9546.
[67] Ehsan Zahedinejad and Arman Zaribafiyan. 2017. Combinatorial Optimization on Gate Model Quantum Computers: A Survey. arXiv:1708.05294 [quant-ph]