



Delft University of Technology

Quantum Data Management From Theory to Opportunities

Hai, Rihan; Hung, Shih Han; Feld, Sebastian

DOI

[10.1109/ICDE60146.2024.00410](https://doi.org/10.1109/ICDE60146.2024.00410)

Publication date

2024

Document Version

Final published version

Published in

Proceedings - 2024 IEEE 40th International Conference on Data Engineering, ICDE 2024

Citation (APA)

Hai, R., Hung, S. H., & Feld, S. (2024). Quantum Data Management: From Theory to Opportunities. In *Proceedings - 2024 IEEE 40th International Conference on Data Engineering, ICDE 2024* (pp. 5376-5381). (Proceedings - International Conference on Data Engineering). IEEE.
<https://doi.org/10.1109/ICDE60146.2024.00410>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Quantum Data Management: From Theory to Opportunities

Rihan Hai
Web Information Systems Group
Delft University of Technology
r.hai@tudelft.nl

Shih-Han Hung
Institute of Information Science
Academia Sinica
shung@cs.utexas.edu

Sebastian Feld
Quantum Machine Learning Group
Delft University of Technology
s.feld@tudelft.nl

Abstract—Quantum computing has emerged as a transformative tool for future data management. Classical problems in database domains, including query optimization, data integration, and transaction management, have recently been addressed using quantum computing techniques. This tutorial aims to establish the theoretical foundation essential for enhancing methodologies and practical implementations in this line of research. Moreover, this tutorial takes a forward-looking approach by delving into recent strides in quantum internet technologies and the nonlocality theory. We aim to shed light on the uncharted territory of future data systems tailored for the quantum internet.

I. INTRODUCTION

In recent years, the realm of computer science has been abuzz with the potential of quantum mechanics. Quantum computing, which leverages the principles of superposition and entanglement, promises computational capacities far beyond what traditional computers can achieve [1]. Moreover, the *quantum internet* holds the potential to offer advantages and capabilities beyond the power of classical internet systems, e.g., secure communication [2] or distributed computing [3]. A quantum internet is a network connecting end nodes that range from simple quantum devices with one qubit to large-scale quantum computers [4]. Recent studies showed the possibility of real-world quantum internet in the scale of kilometers, e.g., 248 kilometers realized using optical fiber [5] and 1203 kilometers with satellite [6]. One vision is that there will be cloud data centers across continents linked by quantum internet, with quantum entanglement promising instantaneous, consistent, and ultimately secure data transmission.

Since the 1960s, database systems have evolved significantly, transitioning from early hierarchical and network models to the widespread adoption of the relational model [7] and relational databases in the 70s-80s. With the birth of the World Wide Web and the surge of web-based applications in the 90s, we witnessed the rise of distributed databases and object-oriented databases [8], [9]. In the 2010s, the emergence of big data technologies, NoSQL databases, and cloud computing further reshaped the landscape of data management [10]. As data continues to grow drastically in both volume and variety, the conventional means of data management will eventually hit the ceiling. Quantum computing, with its revolutionary potential, positions itself as a pivotal future technology in the ongoing evolution of data management systems.





Figure 1: (a) Cloud quantum computer¹ at QuTech, TU Delft. (b) A five-qubit quantum chip using superconducting technology [15]. (c) The basic unit of a quantum internet: end nodes (quantum devices) and a repeater. The repeater establishes quantum entanglement (yellow lines) with each end node, enabling data transmission through quantum teleportation [2].

With the rapid advancement and increasing accessibility of quantum computing technologies [11]–[14], a pressing need arises to understand the intricacies of data management using quantum computers and the quantum internet. Guaranteeing the management and accessibility of data within a quantum computing environment is prominent. This requires understanding and utilizing the innovations of quantum technologies. Even further, it calls for a fundamental rethinking of data handling and management paradigms. *Entering the era of quantum computing, what is the future of data management?*

II. FUNDAMENTALS OF QUANTUM COMPUTING

We will start the tutorial by introducing the basic concepts of quantum computing [1] related to our discussion in Sec. III–IV. For understandability, we introduce below qubits with their analogous concept in classical computation, i.e., bits. We explain qubits at the abstract mathematical level, such that it is general for diverse software and hardware realization.

A. Basics: qubit, superposition, and entanglement

Quantum bit, or qubit, is the fundamental concept of quantum computing. Classical computers use binary digits, bits, as the

¹Quantum Inspire (www.quantum-inspire.com). Photo taken by Marieke de Lorijn for QuTech.

basic unit of information, which can exist in one of two states: 0 or 1. Similarly, two basic states for a qubit are $|0\rangle$ and $|1\rangle$. To represent the state of a quantum system, the term *ket* is used, denoted with the Dirac notation $|\rangle$.

Superposition. The first difference between a classical and quantum bit is: a classical bit can only be 0 or 1, like a coin with a head or tail; but a qubit can exist in a *superposition* of $|0\rangle$ and $|1\rangle$ at the same time. Mathematically, the superposition of a qubit is represented using the linear combination of basic states $|0\rangle$ and $|1\rangle$:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$$

Here, $|\psi\rangle$ is the state of the qubit with α and β being complex numbers. $\alpha|0\rangle + \beta|1\rangle$ means that a qubit would be measured as 0 with the probability of $|\alpha|^2$ and 1 with the probability of $|\beta|^2$.

Example II.1. Consider the qubit in the following superposition:

$$|\psi\rangle = \frac{1}{\sqrt{2}}|0\rangle + \frac{1}{\sqrt{2}}|1\rangle$$

When measuring this qubit, there is an equal probability of 50% ($|\alpha|^2 = |\beta|^2 = \frac{1}{2}$) to get a 0 or 1 as a result.

Entanglement. Another difference between qubits and bits is *entanglement* [16], [17], referred to as the *spooky action at a distance* by Einstein.

Imagine two qubits with qubit A being in Amsterdam and qubit B in San Francisco. When qubit A and qubit B are entangled this means their states are correlated. The *spooky* part is that if we measure qubit A in Amsterdam and find it in a state of 0 (or 1), we instantaneously know the state of qubit B is 0 (or 1), even though it is far away in San Francisco. This correlation is maintained independently of the physical distance between them.

We can consider entangled qubits as a valuable resource in quantum computing [1], which cannot be provided by classical bits. This shared state allows quantum computers to process information in ways that classical computers cannot. A quantum computer with several entangled qubits in superposition can be processed in a single operation, providing a significant advantage for certain problems, e.g., factoring large numbers [18] or searching unsorted databases [19].

B. Quantum computers and libraries

Architecture. Quantum computers come in various models, using different approaches to maintain and manipulate the quantum state of qubits, and are suitable for specific tasks. The most important types of quantum computers include quantum annealers and physical systems for implementing gate-based models such as trapped-ion quantum computers, superconducting qubits, topological quantum computers, and photonic quantum computers. In the tutorial, we will cover both quantum annealers and gate-based models, which are utilized in existing works (cf. Table I). Quantum annealers are designed to solve optimization problems by finding the lowest energy state of a system, which should correspond to

the optimal solution. Quantum annealers can currently be built with more qubits than other existing types of quantum computers, e.g., 5000-qubit D-Wave Advantage quantum annealing processor [32]. Gate-based models are universal, i.e., they are in theory capable of performing any computational task that can be executed on a classical computer, but potentially more efficient for certain problems.

Open-source libraries. For developing applications on quantum computers, key resources include open-source toolkits like IBM's Qiskit², Microsoft's QDK³, and Python libraries such as Google's Cirq⁴, Amazon's OQpy⁵, TU Delft's OpenQL⁶, and Qibo⁷.

III. DATA MANAGEMENT USING QUANTUM COMPUTERS

When it comes to quantum computing applied to data management, we are still in the nascent stages. Pioneering efforts have been made in both the theoretical design and early prototyping of database problems for quantum computers, which we briefly explain in Sec. III-A.

In Table I, more recent works cover essential data management topics such as query optimization [20]–[27], data integration [28], and transaction management [29]–[31]. The general methodology is to map a data management problem to a mathematical formulation solvable by a quantum computer. These works, despite solving different problems each, are mostly mapped a so-called *quadratic unconstrained binary optimization* (QUBO) problem. For quantum computing, currently, QUBO is one of the most widely applied optimization models [33]. During the tutorial, we will elaborate on the existing attempts summarized in Table I. In Sec. III-B, we discuss the query optimization problem, which is one of the cornerstones of database research.

A. Quantum database search and query complexity

A *database* is a collection of data stored electronically in a computer system [34], [35]. In the context of data management, we discuss data complexity, and combined complexity of evaluating database queries [36], [37], e.g., conjunctive queries, in terms of the sizes of input database instances and queries.

In the context of quantum computing, a “quantum database” is a conceptual framework for processing and searching data using quantum algorithms. These algorithms leverage quantum mechanics to address specific problems more efficiently than classical approaches. In this context, a key metric for comparing the computational efficiency of quantum and classical algorithms is *query complexity* [38]. Consider a database with $N = 2^n$ records, in which each record is identified by an n -bit integer label (e.g., “010” for a 3-bit label). The goal is to find a

²<https://qiskit.org/>

³<https://learn.microsoft.com/en-us/azure/quantum/overview-what-is-qsharp-and-qdk>

⁴<https://github.com/quantumlib/Cirq>

⁵<https://github.com/openqasm/oqpy>

⁶<https://openql.readthedocs.io/en/latest/>

⁷<https://github.com/qiboteam/qibo>

Table I: Recent data management works using quantum computers: an overview

Reference	DB problem	Subproblem	Formulation	Intermediate quantum algorithm	Quantum computer
[20]	Query optimization	Multiple query optimization (MQO)	QUBO	–	Annealing-based
[21], [22]				QAOA	Gate-based
[23]–[25]		Join ordering (JO)		QAOA	Gate-based & annealing-based
[26]				QAOA, VQE	Gate-based & annealing-based
[27]				VQC	Gate-based
[28]	Data integration	Schema matching	QUBO	QAOA	Gate-based & annealing-based
[29]–[31]	Transaction management	Two-phase locking	QUBO	–	Annealing-based

target record x_0 using a binary function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, where

$$f(x) = \begin{cases} 1, & \text{if } x = x_0 \\ 0, & \text{otherwise} \end{cases}$$

A quantum algorithm's efficiency can be measured by its query complexity, i.e., how many times it queries this function $f(x)$. One of the most well-known algorithms demonstrating quantum speedup is *Grover's algorithm* [19]. To search a specific record in an unsorted database of N records, classical algorithms require $O(N)$ operations, while Grover's algorithm achieves this in $O(\sqrt{N})$ operations. Grover's algorithm has inspired early research into quantum algorithms for searching databases [39]–[44]. Moreover, efforts have been made to develop quantum query languages akin to SQL for classical databases, covering basic operations: join and set operations (intersection, union, difference) [45]–[50], and data manipulation operations (insert, update, delete) [46], [49], [51].

B. Query optimization

Recent quantum-related works mainly tackled two subproblems of query optimization: multiple query optimization (MQO) and join ordering (JO).

Multiple query optimization studies how to choose query plans given a set of queries [52], which is NP-hard. As one of the earliest attempts to study the MQO problem with modern quantum computers, Trummer et al. proposed in 2016 a method using the D-Wave 2X annealing-based quantum computer [20]. The proposed method includes logical and physical mapping. At the logical level, it transforms the MQO problem to a QUBO formulation, guaranteeing each query has one query plan selected while minimizing the time-wise efficiency. At the physical level, the logical formula of QUBO is further transformed to the energy formula coherent with the physical design of the quantum computer [53]. The experiments demonstrated 1000x speedup on quantum annealer compared to state-of-the-art MQO solutions at that time, although only for a limited subset of MQO problems. Annealing-based quantum computers can reach a larger number of qubits. However, the gate-based model is a more general architecture. Subsequent works [21], [22] tackled the MQO problem on gate-based models using the Quantum Approximate Optimization Algorithm (QAOA) [54], which is a hybrid quantum-classical algorithm.

Join ordering studies how to identify the optimal ordering of join operations between relations for an efficient query plan [55]–[57]. To find optimal left-deep join trees, Schönberger




Figure 2: Roadmap for solving data management problems on quantum computers.

et al. [23], [24] defined the intermediate formulation based on mixed integer linear programming (MILP), followed by binary integer linear programming (BILP) models, which are eventually transformed to QUBO. Recent works also expanded the problem to bushy join trees [25], [26]. In [27], Winker et al. provided a different angle: treating join ordering as a *learning* problem, then apply quantum machine learning techniques based on a variational quantum circuit (VQC) algorithm [58].

C. Opportunities for data management using quantum computers

We outline a few near-term research directions in the following.

1) *Database problem reformulation.* Illustrated in Fig. 2, the general methodology to solve data management problems on quantum computers includes two steps. First, we need to reformulate the data management problem (e.g., MQO or schema matching) to a problem which can be solved on quantum computers (e.g., QUBO). Besides QAOA, there are many more algorithms for quantum computers, e.g., variational quantum eigensolver (VQE), quantum phase estimation (QPE), and Grover's algorithm. Therefore, the main research opportunities are to explore *which* data management problems can be formalized into problems solvable by quantum computers and provide better or faster solutions compared to classical solutions.

2) *Hybrid approach using classical and quantum computers.* Most works in Table I require additional steps on classical computers. For instance, in [20] a prior step is to cluster the queries sharing intermediate results, which significantly reduces the required number of qubits. Many research challenges emerge regarding how to design a hybrid approach using classical and quantum computers.

3) *Optimization given quantum computer constraints.* Despite the compelling computation power of quantum computers, to design data management systems that make

use of quantum computers, we still face many practical constraints such as the restricted number of qubits as well as noisy operations [59]. For a foreseeable period of time, such limitations will coexist while developing data management systems that utilize quantum computers. It poses technical challenges and research opportunities to achieve optimal performance (e.g., time-wise efficiency, accuracy, or completeness of the results), given the physical properties and constraints of quantum computers.

IV. OPPORTUNITIES: DATA MANAGEMENT VIA QUANTUM INTERNET

Quantum internet promises unprecedented new capabilities beyond our internet today [60]. The applications include distributed systems [4], quantum cloud computing [61], quantum key distribution [62]. One of the most fundamental theories enabling these new capabilities is *quantum nonlocality* [63]. Intuitively, we can understand quantum nonlocality as: in a distributed system, we have two computers in different locations; when we change a file on one computer, the corresponding file on the other computer would change immediately. It bypasses our usual assumptions about network distance and speed limitations to transfer data.

The study of quantum nonlocality is at subatomic level, i.e., entangled particles. During the tutorial, we will explain representative nonlocality problems (e.g., CHSH [64] and GHZ [65] games) using logical operators, e.g., AND (\wedge), OR (\vee). This shares a unified mathematical foundation with data management problems. For example, for query optimization, a query plan can be expressed as relation predicates connected by logical operators [66].

A. Quantum nonlocality

In Sec. II, we mentioned superposition and entanglement. To explain the phenomenon of nonlocality, we introduce one of the most well-known states for entanglement, *Bell's state*.

Example IV.1. (Bell's state) We extend the one-qubit state in Example II.1 to a two-qubit state as follows:

$$|\Psi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$$

Nonlocal games. To make it simple, we consider only two players, Alice and Bob. A two-player nonlocal game is a game that consists of two players and a referee. During the game, each player receives an input (like a question or a number) from the referee. Two players, Alice and Bob, however, cannot communicate once the game starts. The intriguing aspect of nonlocal games is that if Alice and Bob share an entangled quantum state, e.g., Bell's state from Example IV.1, they sometimes can win the game with a higher probability than classical strategies. We give an example:

Example IV.2. (Clauser-Horne-Shimony-Holt (CHSH) game [64]) In this game, Alice receives input x and outputs one-bit answers a , while Bob receives y and outputs b . The players win the game if the questions and the answers satisfy

$x \wedge y = a \oplus b$, where \wedge is the logical AND operator and \oplus is XOR. We omit the calculations here, but show results. The two players win optimally with score ≈ 0.85 using an entangled Bell's state, and every pair of players who do not share entangled states can succeed with probability of at most 0.75.

The CHSH game demonstrates an example of entanglement winning over classical strategies by a higher success probability (0.85 vs. 0.75). During the tutorial, we will give another example, the three-player Greenberger-Horne-Zeilinger (GHZ) game [65]. In the GHZ game, the entangled state achieves a probability of 1, while classical resources can only achieve a probability of 0.75. That is, with entanglement, we can achieve a task that is *not* possible with classical resources.

B. Future research directions of data management via quantum internet

Quantum nonlocality serves as the theoretical foundation of protocols for secure communication and key distribution [62]. Naturally, it is an interesting direction to explore secure data management via quantum internet. Next, we discuss new design challenges and research opportunities, which are not discussed in the existing literature.

1) *New data structures.* We face the fundamental challenge of data representation, since communicating quantum information is fundamentally different from classical information. Take the famous *no-cloning theorem* [1] for example. Existing database theory and systems are built upon the fact that we can freely read and copy data, e.g., duplicate a data tuple, or replicate a dataset (e.g., a CSV file) on another computer in the network. However, the no-cloning theorem states that we cannot make an exact copy of an arbitrary quantum state. Thus: How to design data models, when quantum data cannot be copied without destroying the original version?

2) *New architectures of distributed data management systems.* Distributed database systems traditionally manage distributed, fragmented, and replicated data over classical computer networks. Facing quantum properties such as entanglement, nonlocality, and no-cloning theorem, we need new system architectures to ensure reliability, availability, and scalability. The question arises, for instance, how to design distributed data management systems based on protocols used in quantum internet [4]? Given the quantum hardware or network failures [67], how to design fault tolerance and recovery strategies?

V. PRESENTERS

Rihan Hai is an assistant professor at TU Delft, The Netherlands. Her research focuses on data lakes, data management for machine learning, and quantum data management.

Shih-Han Hung is a postdoc at Academia Sinica. His research aims to better understand the power and the limit of quantum computers.

Sebastian Feld is an assistant professor at TU Delft, The Netherlands. He and his group are working on Quantum Machine Learning.

REFERENCES

- [1] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information*. Cambridge university press, 2010.
- [2] A. Broadbent, J. Fitzsimons, and E. Kashefi, "Universal blind quantum computation," in *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, 2009, pp. 517–526.
- [3] N. H. Nickerson, J. F. Fitzsimons, and S. C. Benjamin, "Freely scalable quantum technologies using cells of 5-to-50 qubits with very lossy and noisy photonic links," *Physical Review X*, vol. 4, no. 4, p. 041041, 2014.
- [4] S. Wehner, D. Elkouss, and R. Hanson, "Quantum internet: A vision for the road ahead," *Science*, vol. 362, no. 6412, p. eaam9288, 2018.
- [5] S. P. Neumann, A. Buchner, L. Bulla, M. Bohmann, and R. Ursin, "Continuous entanglement distribution over a transnational 248 km fiber link," *Nature Communications*, vol. 13, no. 1, p. 6134, 2022.
- [6] J. Yin, Y. Cao, Y.-H. Li, S.-K. Liao, L. Zhang, J.-G. Ren, W.-Q. Cai, W.-Y. Liu, B. Li, H. Dai *et al.*, "Satellite-based entanglement distribution over 1200 kilometers," *Science*, vol. 356, no. 6343, pp. 1140–1144, 2017.
- [7] E. F. Codd, "A relational model of data for large shared data banks," *Communications of the ACM*, vol. 13, no. 6, pp. 377–387, 1970.
- [8] M. Qszo and P. Valdureiz, "Principles of distributed database system," 1991.
- [9] M. Atkinson, D. Dewitt, D. Maier, F. Bancilhon, K. Dittrich, and S. Zdonik, "The object-oriented database system manifesto," in *Deductive and object-oriented databases*. Elsevier, 1990, pp. 223–240.
- [10] D. Abadi, R. Agrawal, A. Ailamaki, M. Balazinska, P. A. Bernstein, M. J. Carey, S. Chaudhuri, J. Dean, A. Doan, M. J. Franklin *et al.*, "The beckman report on database research," *Communications of the ACM*, vol. 59, no. 2, pp. 92–99, 2016.
- [11] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell *et al.*, "Quantum supremacy using a programmable superconducting processor," *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.
- [12] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu *et al.*, "Quantum computational advantage using photons," *Science*, vol. 370, no. 6523, pp. 1460–1463, 2020.
- [13] Y. Wu, W.-S. Bao, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan *et al.*, "Strong quantum computational advantage using a superconducting quantum processor," *Physical review letters*, vol. 127, no. 18, p. 180501, 2021.
- [14] L. S. Madsen, F. Laudenbach, M. F. Askarani, F. Rortais, T. Vincent, J. F. Bulmer, F. M. Miatto, L. Neuhaus, L. G. Helt, M. J. Collins *et al.*, "Quantum computational advantage with a programmable photonic processor," *Nature*, vol. 606, no. 7912, pp. 75–81, 2022.
- [15] R. van der Heijden, "The first european cloud quantum computer is in delft," <https://www.nemokenislink.nl/publicaties/eerste-europese-cloudquantumcomputer-staat-in-delft>, accessed: 2024-02-23.
- [16] A. Einstein, B. Podolsky, and N. Rosen, "Can quantum-mechanical description of physical reality be considered complete?" *Physical review*, vol. 47, no. 10, p. 777, 1935.
- [17] E. Schrödinger, "Discussion of probability relations between separated systems," in *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 31, no. 4. Cambridge University Press, 1935, pp. 555–563.
- [18] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM review*, vol. 41, no. 2, pp. 303–332, 1999.
- [19] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, 1996, pp. 212–219.
- [20] I. Trummer and C. Koch, "Multiple query optimization on the D-Wave 2X adiabatic quantum computer," *Proceedings of the VLDB Endowment*, vol. 9, no. 9, 2016.
- [21] T. Fankhauser, M. E. Solèr, R. M. Füchslin, and K. Stockinger, "Multiple query optimization using a hybrid approach of classical and quantum computing," *arXiv preprint arXiv:2107.10508*, 2021.
- [22] T. Fankhauser, M. E. Soler, R. M. Füchslin, and K. Stockinger, "Multiple query optimization using a gate-based quantum computer," *IEEE Access*, 2023.
- [23] M. Schönberger, "Applicability of Quantum Computing on Database Query Optimization," in *Proceedings of the 2022 International Conference on Management of Data*, ser. SIGMOD '22. New York, NY, USA: Association for Computing Machinery, 2022, pp. 2512–2514.
- [24] M. Schönberger, S. Scherzinger, and W. Maurer, "Ready to leap (by co-design)? join order optimisation on quantum hardware," *Proceedings of the ACM on Management of Data*, vol. 1, no. 1, pp. 1–27, 2023.
- [25] M. Schönberger, I. Trummer, and W. Maurer, "Quantum Optimisation of General Join Trees," in *Joint Workshops at 49th International Conference on Very Large Data Bases (VLDBW'23)—International Workshop on Quantum Data Science and Management (QDSM'23)*, August 28–September 1, 2023, Vancouver, Canada (CEUR Workshop Proceedings), 2023.
- [26] N. Nayak, J. Rehfeld, T. Winker, B. Warnke, U. Çalikylmaz, and S. Groppe, "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*, ser. BiDEDE '23. New York, NY, USA: Association for Computing Machinery, 2023.
- [27] T. Winker, U. Çalikylmaz, L. Gruenwald, and S. Groppe, "Quantum Machine Learning for Join Order Optimization Using Variational Quantum Circuits," in *Proceedings of the International Workshop on Big Data in Emergent Distributed Environments*, ser. BiDEDE '23. New York, NY, USA: Association for Computing Machinery, 2023.
- [28] K. Fritsch and S. Scherzinger, "Solving hard variants of database schema matching on quantum computers," *Proceedings of the VLDB Endowment*, vol. 16, no. 12, pp. 3990–3993, 2023.
- [29] T. Bittner and S. Groppe, "Avoiding blocking by scheduling transactions using quantum annealing," in *Proceedings of the 24th Symposium on International Database Engineering and Applications*, 2020.
- [30] Tim Bittner and Sven Groppe, "Hardware accelerating the optimization of transaction schedules via quantum annealing by avoiding blocking," *Open Journal of Cloud Computing (OJCC)*, vol. 7, no. 1, pp. 1–21, 2020.
- [31] S. Groppe and J. Groppe, "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*, 2021, pp. 149–156.
- [32] A. D. King, J. Raymond, T. Lanting, R. Harris, A. Zucca, F. Altomare, A. J. Berkley, K. Boothby, S. Ejtemaee, C. Enderud *et al.*, "Quantum critical dynamics in a 5,000-qubit programmable spin glass," *Nature*, pp. 1–6, 2023.
- [33] F. Glover, G. Kochenberger, and Y. Du, "A tutorial on formulating and using qubo models," *arXiv preprint arXiv:1811.11538*, 2018.
- [34] S. Abiteboul, R. Hull, and V. Vianu, *Foundations of databases*. Addison-Wesley Reading, 1995, vol. 8.
- [35] R. Ramakrishnan, J. Gehrke, and J. Gehrke, *Database management systems*. McGraw-Hill New York, 2003, vol. 3.
- [36] M. Y. Vardi, "The complexity of relational query languages," in *Proceedings of the fourteenth annual ACM symposium on Theory of computing*, 1982, pp. 137–146.
- [37] C. H. Papadimitriou and M. Yannakakis, "On the complexity of database queries," in *Proceedings of the sixteenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, 1997, pp. 12–19.
- [38] A. Ambainis, "Understanding quantum algorithms via query complexity," in *Proceedings of the International Congress of Mathematicians: Rio de Janeiro 2018*. World Scientific, 2018, pp. 3265–3285.
- [39] B. M. Terhal and J. A. Smolin, "Single quantum querying of a database," *Physical Review A*, vol. 58, no. 3, p. 1822, 1998.
- [40] M. Boyer, G. Brassard, P. Høyer, and A. Tapp, "Tight bounds on quantum searching," *Fortschritte der Physik: Progress of Physics*, vol. 46, no. 4–5, pp. 493–505, 1998.
- [41] A. Patel, "Quantum database search can do without sorting," *Physical Review A*, vol. 64, no. 3, p. 34303, 2001.
- [42] I.-M. Tsai, S.-Y. Kuo, and D. S. L. Wei, "Quantum Boolean circuit approach for searching an unordered database," in *Proceedings of the 2nd IEEE Conference on Nanotechnology*. IEEE, 2002, pp. 315–318.
- [43] S. Imre and F. Balázs, "The generalized quantum database search algorithm," *Computing*, vol. 73, pp. 245–269, 2004.
- [44] Y.-L. Ju, I.-M. Tsai, and S.-Y. Kuo, "Quantum circuit design and analysis for database search applications," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 54, no. 11, pp. 2552–2563, 2007.
- [45] P. Cockshott, "Quantum relational databases," *arXiv preprint quant-ph/9712025*, 1997.
- [46] A. Gueddana, R. Chatta, and N. Boudriga, "Optimized methods for inserting and deleting records and data retrieving in quantum database," in *2010 12th International Conference on Transparent Optical Networks*. IEEE, 2010, pp. 1–5.

- [47] T. Salman and Y. Baram, "Quantum set intersection and its application to associative memory," *The Journal of Machine Learning Research*, vol. 13, no. 1, pp. 3177–3206, 2012.
- [48] C.-Y. Pang, R.-G. Zhou, C.-B. Ding, and B.-Q. Hu, "Quantum search algorithm for set operation," *Quantum information processing*, vol. 12, no. 1, pp. 481–492, 2013.
- [49] A. Gueddana, R. Chatta, and M. Attia, "CNOT-based design and query management in quantum relational databases," *International Journal of Quantum Information*, vol. 12, no. 04, p. 1450023, 2014.
- [50] S. Józczik and A. Kiss, "Quantum Computation and Its Effects in Database Systems," in *European Conference on Advances in Databases and Information Systems*. Springer, 2020, pp. 13–23.
- [51] A. Younes, "Database Manipulation Operations on Quantum Systems," *Quantum Information Review*, vol. 1, pp. 9–17, 2013.
- [52] T. K. Sellis, "Multiple-query optimization," *ACM Transactions on Database Systems (TODS)*, vol. 13, no. 1, pp. 23–52, 1988.
- [53] C. C. McGeoch and C. Wang, "Experimental evaluation of an adiabatic quantum system for combinatorial optimization," in *Proceedings of the ACM International Conference on Computing Frontiers*, 2013, pp. 1–11.
- [54] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm," *arXiv preprint arXiv:1411.4028*, 2014.
- [55] M. Steinbrunn, G. Moerkotte, and A. Kemper, "Heuristic and randomized optimization for the join ordering problem," *The VLDB journal*, vol. 6, pp. 191–208, 1997.
- [56] V. Leis, A. Gubichev, A. Mirchev, P. A. Boncz, A. Kemper, and T. Neumann, "How good are query optimizers, really?" *Proc. VLDB Endow.*, vol. 9, no. 3, pp. 204–215, 2015. [Online]. Available: <http://www.vldb.org/pvldb/vol9/p204-leis.pdf>
- [57] I. Trummer and C. Koch, "Solving the join ordering problem via mixed integer linear programming," in *Proceedings of the 2017 ACM International Conference on Management of Data*, 2017, pp. 1025–1040.
- [58] S. Y.-C. Chen, C.-H. H. Yang, J. Qi, P.-Y. Chen, X. Ma, and H.-S. Goan, "Variational quantum circuits for deep reinforcement learning," *IEEE Access*, vol. 8, pp. 141 007–141 024, 2020.
- [59] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf, "Introduction to quantum noise, measurement, and amplification," *Reviews of Modern Physics*, vol. 82, no. 2, p. 1155, 2010.
- [60] W. Kozłowski and S. Wehner, "Towards large-scale quantum networks," in *Proceedings of the sixth annual ACM international conference on nanoscale computing and communication*, 2019, pp. 1–7.
- [61] D. Castelvecchi, "Ibm's quantum cloud computer goes commercial," *Nature*, vol. 543, no. 7644, 2017.
- [62] C. H. Bennett and G. Brassard, "Quantum cryptography: Public key distribution and coin tossing," *Theoretical computer science*, vol. 560, pp. 7–11, 2014.
- [63] H. Buhman, R. Cleve, S. Massar, and R. de Wolf, "Nonlocality and communication complexity," *Rev. Mod. Phys.*, vol. 82, pp. 665–698, Mar 2010. [Online]. Available: <https://link.aps.org/doi/10.1103/RevModPhys.82.665>
- [64] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, "Proposed experiment to test local hidden-variable theories," *Physical review letters*, vol. 23, no. 15, p. 880, 1969.
- [65] D. M. Greenberger, M. A. Horne, and A. Zeilinger, "Going beyond bell's theorem," in *Bell's theorem, quantum theory and conceptions of the universe*. Springer, 1989, pp. 69–72.
- [66] F. Kastrati and G. Moerkotte, "Generating optimal plans for boolean expressions," in *2018 IEEE 34th International Conference on Data Engineering (ICDE)*. IEEE, 2018, pp. 1013–1024.
- [67] Y. Dulek, A. B. Grilo, S. Jeffery, C. Majenz, and C. Schaffner, "Secure multi-party quantum computation with a dishonest majority," in *Annual International Conference on the Theory and Applications of Cryptographic Techniques*. Springer, 2020, pp. 729–758.