

A Toolbox to Understand the Physics of Quantum Data Management

Wolfgang Mauerer

Technical University of Applied Sciences Regensburg
Siemens AG, Foundational Technology
Regensburg/Munich, Germany

Manuel Schönberger

Cornell University
Ithaca, NY, USA

Abstract

The application of quantum computing to data management has attracted growing interest, yet remains constrained by a limited understanding of how the physical behaviour of quantum devices relates to the structure and difficulty of database problems. In particular, evaluating quantum annealing approaches for combinatorial optimisation, which is central to many data management tasks, poses significant challenges beyond the scope of conventional empirical and complexity-theoretic methods.

We present a computational toolbox for the systematic numerical analysis of quantum annealing processes derived from data management problem formulations. Adopting a physics-informed perspective, the toolbox enables the study of spectral and dynamical properties – such as energy gaps and eigenstate structure – that are inaccessible through direct hardware measurements, yet essential for understanding computational hardness and scaling behaviour.

Our approach further provides derived quantities and visualisation techniques that support the interpretation of optimisation dynamics, the identification of structural similarities to canonical physical models, and the construction of reduced effective descriptions. By bridging methodological gaps between quantum computing and database systems research, this work establishes a principled foundation for evaluating quantum approaches and guiding future co-design efforts.

CCS Concepts

• **Theory of computation** → *Discrete optimization; Database query processing and optimization (theory)*; • **Computer systems organization** → **Quantum computing**; • **Computing methodologies** → **Quantum mechanic simulation**.

Keywords

Quantum Annealing, Quantum DBMS, Multi-Query Optimisation, Empirical analysis

ACM Reference Format:

Wolfgang Mauerer and Manuel Schönberger. 2026. A Toolbox to Understand the Physics of Quantum Data Management. In *The third workshop on Quantum Computing and Quantum-Inspired Technology for Data-Intensive Systems and Applications (Q-Data '26)*, May 31–June 05, 2026, Bengaluru, India. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3811628.3811835>



This work is licensed under a Creative Commons Attribution 4.0 International License. *Q-Data '26, Bengaluru, India*

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2703-0/2026/05
<https://doi.org/10.1145/3811628.3811835>

1 Introduction

The use of quantum computing to accelerate or otherwise improve over classical baselines has seen a growing interest in the last years [6, 10, 13, 67]. While limitations of quantum computers, especially their inability to process even moderately large amounts of classical data [20, 28], clearly limit the possible use-cases in DBMS, many fundamental and long-standing problems in data management concern the solution of combinatorial optimisation problems that work on large conceptual search spaces without the need for large concrete data. Discussions about scaling of and advantage from quantum combinatorial optimisation arise plentiful in the literature [7, 8, 11, 14, 15, 17, 18, 23, 24, 24, 31, 31, 36, 45, 56–58, 62, 63, 65] (see also Sec. 5). Predicting and understanding the performance of quantum approaches is a generic problem [37]: An empirical analysis is mostly impossible given the unavailability of large enough quantum hardware and the noisiness of current systems. A theoretical analysis requires an entirely different approach than for classical algorithms, and is computationally challenging and typically as hard as solving the subject problem in the first place.

In this paper, we present a tool for the classical numerical analysis of quantum annealing problems, and introduce, from a physics-informed perspective, a judiciously chosen set of quantities and visualisations that can be used to further the understanding of links between the complexity-theoretic properties of such problems, and the physical properties of quantum annealers that are employed to solve them. A general overview about our tool, available as fully reproducible [39] open source software on <https://github.com/lfd/anneal>, is provided in Fig. 1.

While the theory behind many aspects of DBMS is extremely well developed, especially practical systems rely on a plethora of stochastic algorithms, heuristics or machine learning approaches. Any suggested quantum approach needs to be evaluated against such baselines, which is not possible empirically owing to the limitations of current prototypical machines. As the analysis of computational complexity (and other properties) of quantum annealing strongly differs from what is established in computer science, but is usually also impossible via a simple input-output relationship between problem and characteristics, new means of judging and predicting potential or expected performance of quantum approaches to DBMS problem are required.

Our goal is to establish a computational pipeline that allows researchers to address, among others, the following: (a) understand dynamical patterns and computational hardness of quantum approaches used for data management (and, more general, combinatorial optimisation) problems; (b) compute quantitative and theoretically grounded estimates on scaling, success probability, etc.

for such formulations, and enable performance predictions outside the empirically accessible problem range; (c) reveal similarities of problem formulations in data management with seminal computational problems, also from the physics literature; (d) construct surrogate simplified models based on effective dynamics that can be used to construct quantum-inspired approaches, or gain a deeper understanding of the properties of the original problem; (e) avoid interpretation traps and pitfalls based on an incomplete or superficial understanding of the underlying physical computational processes; and (f) identify possibilities for co-design approaches that allow for designing quantum accelerators with a focus on data management challenges

Based on our experience, we observe a certain amount of disconnect between established practices for theoretically and empirically assessing the guarantees and potential advantages of quantum approaches, and the expectations commonly applied in the evaluation of database systems research. In particular, methods that are standard and well-justified within the quantum computing literature do not always align with conventional benchmarking and evaluation paradigms in the DBMS community. Work that follows established evaluation methodologies that may, for quantum approaches, not always be meaningful or appropriate, may be more readily accepted, whereas approaches that employ technically sound yet less familiar evaluation techniques can face additional scrutiny. One of the goals of this paper is therefore to raise awareness of this mismatch and to encourage a more nuanced discussion of evaluation criteria at the intersection of quantum computing and data management.

To judge hardness and scaling behaviour in quantum optimisation requires considerably more knowledge than can be obtained from simple empirical measurements of success probabilities and quality of approximation measurements. For instance, order parameters and their relation to occurring phase transitions are seen as crucial in the physics-centric literature, but rarely if ever employed for DBMS. Yet, the order of the underlying phase transition is a property of the spectrum and ground-state structure of $H(s)$, not of one particular finite-rate run [46], which make more advanced and general empirical analysis techniques necessary. Our tool provides the necessary ingredients to conduct such analyses.

Even if the general overarching problem – using (future) quantum computers to solve combinatorial optimisation problems – is not uniquely related to problems in data management, we argue that a co-design approach [35, 50, 56], as well as careful tailoring of different stages of the computational and conceptual pipeline from problem formulation to systems integration is necessary to not only optimise benefits, but possibly also required to reach any form of practical quantum advantage at all: The potentials of optimising at the level of problem representation and transformation was, incidentally, one of the early considerations in quantum data management [57, 63], but is increasingly addressed via automatic tools [54, 55], with often considerable differences between algebraically identical formulations. We expect that a better understanding of the underlying computational processes can provide crucial input to such efforts.

While a considerable body of work addresses the analysis of the annealing process as motivated by device physics, relating problems in data management to this point of view is highly non-trivial and not yet sufficiently established. It is useful to be able to control the

difficulty of constructed instances, and construct instances for arbitrary desired sizes. This is usually not the case for DBMS problems that are not continuous in the size of the generated Hamiltonians, and display no immediate connection to Ising models.

2 Concepts and Foundations

As we discuss an interdisciplinary subject that goes beyond the physical depth typically covered in the literature on quantum approaches for data management problems, we need to establish some foundations and consistent terminology. While it is not possible to provide a self-contained introduction into the physical and mathematical aspects of quantum mechanics, we nonetheless summarise the most essential aspects that are required to understand our approach. We assume familiarity with standard concepts of quantum computing as increasingly adopted by the DBMS community.

2.1 Quantum Annealing and Combinatorial Optimisation

The Hamiltonian \hat{H} is a Hermitian operator that represents the total energy of a given system of interest. This can be any microscopic artefact like an atom, a molecule, or a constructed entity like a spin glass whose properties are well suited to map combinatorial optimisation problems onto. The latter is the basis for methods like quantum annealing that have received growing interest in the data management community. While deriving a Hamiltonian for a concrete problem is typically left to physics, it is important to realise that it plays two key roles: Firstly, the eigenvalues E_i of \hat{H} are the possible discrete energy levels¹ that a given system can assume, each of which is accompanied by an eigenstate $|E_i\rangle$, with $\hat{H}|E_i\rangle = E_i|E_i\rangle$. The eigenvectors comprise an orthogonal set. Secondly, the time evolution of a quantum system is governed by the time-dependent Schrödinger equation

$$i\hbar \frac{d}{dt} |\psi(s(t))\rangle = H(s(t)) |\psi(s(t))\rangle, \quad (1)$$

that contains the itself time-varying Hamiltonian $\hat{H}(s(t))$ (we use a normalised time $s(t)$ defined by a monotonically increasing function s that runs from 0 to 1, and defines the annealing schedule) as a central ingredient. Knowledge of these two properties is essential to understand the behaviour of a system governed by a Hamiltonian, and therefore also *key to understanding the properties and behaviour of computational approaches based on Hamiltonian dynamics*, particularly quantum annealing, but also related efforts like QAOA. Obtaining these quantities to understand the effective computational process is our core goal.

Consider that at each value of s , the current state of the quantum annealer can be expanded in the *instantaneous eigenbasis*

$$|\psi(s(t))\rangle = \sum_i c_i(t) |E_i(s(t))\rangle. \quad (2)$$

Inserting this expansion into the Schrödinger equation (essentially following textbook quantum mechanics, e.g. Ref. [42]) yields coupled differential equations for the amplitudes $c_i(t)$, where transitions between eigenstates are governed by matrix elements of $\frac{d}{ds} H(s)$ and inversely by energy differences $E_j(s) - E_i(s)$. Note that

¹Hermiticity of \hat{H} guarantees a real eigen-spectrum, and thus real energy values.

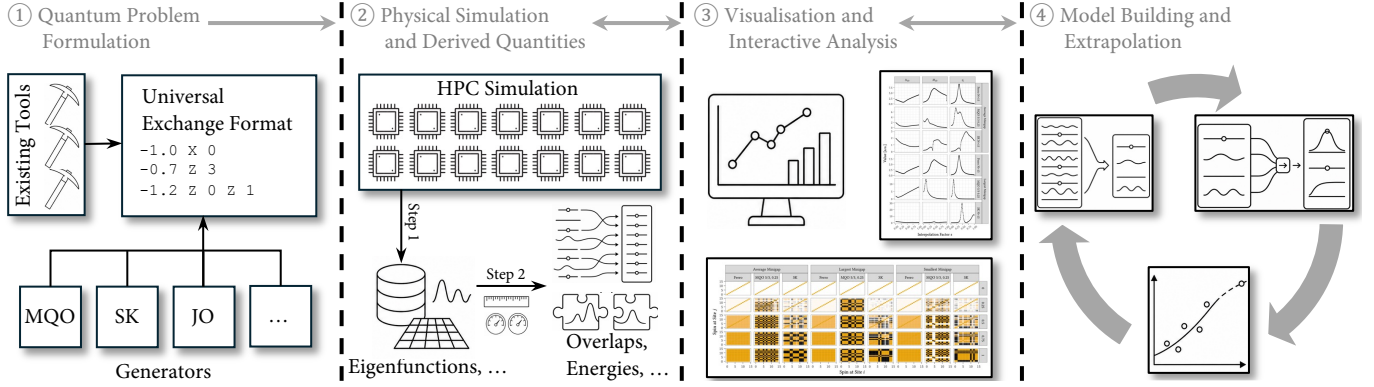


Figure 1: Overview of our Toolbox: Based on existing quantum formulations for optimisation problems that can be transcribed into a simple intermediate format, we use HPC systems to simulate the dynamics of a quantum annealing process. After post-processing the results (deliberately decoupled from the long-running and compute intensive simulation) to obtain quantities of interest, an accompanying set of visualisation scripts is used to understand the salient characteristics of the problem, and the implications on performance, achievable result quality, etc.; this can lead to (simplified) models, surrogates and others that allow for characterising properties of approaches, or predict characteristics.

energies and eigenstates change when the Hamiltonian changes, and re-computing the quantities is necessary.

This is particularly relevant for the class of problems arising in *quantum annealing*, where the goal is to prepare the ground state (with the smallest possible energy E_0 of a problem-specific Hamiltonian \hat{H}_P). This is, depending on the characteristics of \hat{H}_P , known to be a very hard computational problem [16]. Nonetheless, it can be solved by initialising the system in the ground state of a *simple driver Hamiltonian* H_I with an easily prepared ground state, and then evolving the system according to an interpolation (a time-varying Hamiltonian)

$$\hat{H}(s) = A(s)\hat{H}_I + B(s)\hat{H}_P, \quad s \in [0, 1], \quad (3)$$

where H_P encodes the target optimisation problem (in the standard annealing protocol, $\hat{H}_I = -\Gamma \sum_{i=1}^n \hat{\sigma}_x^i$ is used, with a known minimum-energy state $|+\rangle^{\otimes n}$). Quantum methods are often equated with quadratic unconstrained binary optimisation (QUBO) problems. More general choices are beyond our scope, but the link between QUBOs and \hat{H}_P is straightforward: A QUBO is a classical cost function $C(x) = \sum_i a_i x_i + \sum_{i < j} b_{ij} x_i x_j$ with $x_i \in \{0, 1\}$. To embed this into a quantum annealer, binary variables are mapped to spins via $x_i = \frac{1}{2}(1 - \hat{\sigma}_i^z)$, which converts the cost function into a so-called Ising Hamiltonian $\hat{H}_P = \sum_i h_i \hat{\sigma}_i^z + \sum_{i < j} J_{ij} \hat{\sigma}_i^z \hat{\sigma}_j^z$ (we equate spins and binary variables in the following).

This Hamiltonian is diagonal in the computational basis, and each bit-string encoded in a quantum state corresponds to a classical configuration with energy equal (up to a constant shift) to the original QUBO cost. Minimising the QUBO is equivalent to finding the ground state of \hat{H}_P .

Note that the interpolation between initial and final Hamiltonian in Eq. 3 is guided by the annealing parameter s that leads the system from $\hat{H}(0) = \hat{H}_I$ to $\hat{H}(1) = \hat{H}_P$ over an annealing time T . The schedule $A(s), B(s)$ determines how \hat{H}_I changes into \hat{H}_P over time. At $s = 0$, the system is initialised in the ground state $|E_0(0)\rangle$ of \hat{H}_I . The objective of quantum annealing is to evolve the system

towards $s = 1$ such that the final state approximates the ground state $|E_0(s = 1)\rangle$ of \hat{H}_P . As \hat{H}_P is designed such that a minimum-energy eigenstate encodes the solution to a computational problem, the annealing process can therefore be used to find solutions for this problem. In the ideal adiabatic regime, where $s(t)$ varies sufficiently slowly, transitions from the lowest-energy state into higher states are suppressed, and the system follows a single instantaneous eigenstate. In particular, if the system is initialised in $|E_0(0)\rangle$, then

$$|\psi(t)\rangle \approx e^{i\phi(t)} |E_0(t)\rangle, \quad (4)$$

up to a complex phase factor $\phi(t)$ that is inconsequential in a measurement.

Appropriately chosen schedules can not only influence how quickly, but also with which success probability the process reaches a desired final state [1], or which final energy state can be expected with which probability – which in turn influences solution quality (for the rest of this paper, we restrict our considerations to a standard linear schedule $A(s) = (1 - s)$, $B(s) = s$; the tool supports arbitrary schedules though, as their computational handling is not much different from a linear schedule). In contrast to the adiabatic regime, *diabatic* evolution refers to computations that involve transitions between eigenstates; the system explores excited states due to insufficient runtime or small gaps. This leads to computational results observed as incorrect solutions, albeit some algorithms deliberately exploit diabatic evolution [12, 30, 44, 52].

For each fixed s , the Hamiltonian admits the aforementioned spectral decomposition in terms of the time-dependent energy eigenstates²

$$\hat{H}(s) = \sum_i E_i(s) |E_i(s)\rangle \langle E_i(s)|, \quad (5)$$

²While the use of bracket notation is common in computer-science centric applications of quantum computing, the quantity $|E_i\rangle \langle E_i|$ is not ubiquitously used. For a finite-dimensional view of quantum mechanics, which is usually fully sufficient in computer science, $|E_i\rangle$ corresponds to a complex column vector of dimension 2^n , whereas $\langle E_i|$ is the conjugate transpose, that is, a row vector. The outer product $|E_i\rangle \langle E_i|$ of a vector and its dual vector therefore corresponds to a $\mathbb{C}^{2^n \times 2^n}$ matrix that acts as an operator. As it is Hermitian by definition, it can be used as a contribution to a Hamiltonian.

where both eigenvalues and eigenstates depend smoothly on s (under mild conditions). Given knowledge of \hat{H}_I and \hat{H}_P , numerical methods can be used to obtain this decomposition, respectively the involved eigenstates and their eigenvalues. It is important to underline that performing an annealing protocol on real quantum hardware does *not* deliver this information, as usually only probabilities of observing certain final states are recorded. Even more involved measurement schemes do not deliver all information required to obtain the spectral decomposition and the temporal dynamics. Additionally, imperfect hardware [1, 22, 38] as is available today will further perturb any empirical observations, adding another degree of complications into the picture.

Consequently, numerical simulation of the process as performed by the tool delivers *more* information about the dynamics and the computational behaviour as most empirical real-hardware studies.

While ideal adiabatic evolution tracks the ground state, practical implementations require considering a small number of low-energy eigenstates. For each s , define the truncated subspace

$$\mathcal{H}_{\text{eff}}(s) = \text{span}\{|E_0(s)\rangle, \dots, |E_k(s)\rangle\}. \quad (6)$$

Instead of requiring perfect adiabaticity, the state can be approximated as

$$|\psi(s(t))\rangle \approx \sum_{i=0}^k c_i(s(t)) |E_i(s(t))\rangle, \quad (7)$$

with k small. This viewpoint naturally arises when diagonalising $H(s)$ at discrete values of s and propagating the state within the subspace spanned by a limited number of eigenstates, as is at the core of our numerical approach. The validity of this truncation depends on the structure of the spectrum along the interpolation path.

Another central quantity is the instantaneous spectral gap

$$\Delta(s) := \Delta_{10}(s) := E_1(s) - E_0(s), \quad (8)$$

More generally, $\Delta_{k+1,k}(s) := E_{k+1}(s) - E_k(s)$ and $\Delta_{k0}(s) := E_k(s) - E_0(s)$ denote gaps between higher energy levels that are of interest when considering larger subspaces and more involved dynamics, as they appear in real-world quantum annealers. If one of these gaps remains sufficiently large at some energy onwards, transitions to higher-energy states are suppressed, and the dynamics remain effectively confined to the effective Hilbert space $\mathcal{H}_{\text{eff}}(s)$ of the model. Likewise, $\min_s (\Delta_{10}(s))^{-2}$ is known to be a first rough predictor for the required annealing runtime [1].

Let us again highlight that these spectral gap cannot be obtained from concrete measurements on a device without further ado, but are crucially relevant to determine and extrapolate scaling behaviour. This makes any attempt to perform empirical analyses on a concrete device necessarily incomplete.

The restriction to a low-energy subspace along the annealing path is justified when the following conditions hold:

- The system is initialised in the ground state of H_I .
- The interpolation schedule $s(t)$ varies sufficiently slowly relative to inverse spectral gaps.
- Energy gaps separating $\mathcal{H}_{\text{eff}}(s)$ from higher-energy states remain non-negligible.
- Matrix elements inducing transitions to higher-energy states are small.

Under these assumptions, high-energy components remain weakly populated. Moreover, any residual contributions acquire rapidly varying phases and have limited impact on observables, so that the effective dynamics are governed by a small number of instantaneous eigenstates. Conversely, the approximation breaks down near avoided crossings with small gaps, under fast schedules, or in systems with dense spectra. In such cases, transitions out of the low-energy subspace become significant and must be explicitly accounted for.

From this perspective, quantum annealing can be interpreted as a controlled traversal of a sequence of low-dimensional eigenspaces of $H(s)$, where the computational task reduces to maintaining overlap with the evolving ground state (or a small set of low-energy states) of the interpolating Hamiltonian. As energy eigenvalues and eigenstates are of particular importance, careful consideration is required. Recall that at any fixed value of s , the Hamiltonian $H(s)$ has *instantaneous eigenvalues*

$$E_0(s) \leq E_1(s) \leq E_2(s) \leq \dots \quad (9)$$

and corresponding *eigenstates* $|0(s)\rangle, |1(s)\rangle, \dots$. A key subtlety is how the energy levels are labelled as s varies. Two different conventions are commonly discussed: *sorted energy branches* and *tracked energy branches*. These correspond to two different ways of interpreting the spectral data, both supported by the tool.

In the sorted convention, eigenvalues are re-ordered at each s so that $E_0(s)$ is always the lowest eigenvalue, $E_1(s)$ the second lowest, and so on. This ordering is local in s and does not attempt to preserve the identity of the states across different values of s . The adiabatic theory of quantum annealing is formulated in terms of these sorted branches, since the evolution ideally follows the instantaneous ground state $|E_0(s)\rangle$. Consequently, quantities such as the minimum gap $\Delta_{\min} = \min_{s \in [0,1]} \Delta_{10}(s)$ are defined using sorted energy levels.

Eigenstates can alternatively be tracked across different values of s by matching them according to overlap ('similarity') between neighbouring eigenvectors (also illustrated in Fig. 1). This produces *continuous* curves that represent the evolution of a specific physical state as the Hamiltonian changes. These tracked branches can cross each other and therefore do not necessarily remain ordered by energy. They are useful for interpreting how the character of the quantum state changes during the anneal and for identifying tunnelling between configurations.

Near an *avoided crossing*, two tracked branches approach each other closely. In the sorted picture the labels of the levels exchange roles, whereas in the tracked picture the curves cross. This distinction is important: performance metrics such as the minimum gap or adiabatic condition must be computed using the sorted eigenvalues, while tracked branches are primarily used to visualise the physical structure of the transition.

2.1.1 Scalability. The problem we need to solve places considerable computational demands that cannot be provided on typical machines but for the smallest instances (and is also more complex than the widely used simulation of quantum circuits), but requires the use of HPC methods. We rely on established and proven approaches: Scalability is achieved through MPI parallelism that especially distributes huge state vectors across processes; OpenMP threading

that allows for shared-memory parallelism within each process, and most importantly matrix-free evaluation techniques minimises memory bandwidth requirements. We observe good scalability of MPI-level parallelism with up to about 150 cores; this is sufficient to handle one single problem formulation with 25 qubits in roughly one full day. The required time for an analysis with eight eigenpairs and 200 steps of resolution for more moderate system dimensions up to 20 qubits, as well as the storage requirements, are shown in Tab. 1 to familiarise researchers with the expected resource requirements. Beyond around 150 cores, communication and parallelisation overheads eliminate further computational gains, a resorting to parameter-level parallelisation is required. This can be conveniently established for sweeps and parameter scans.

Table 1: Storage requirement $|E|$ for eigenpairs at different system dimensions (# of qubits) N , and for a complete system analysis $|S|$ with eight eigenvectors computed at a resolution of 200 steps. Ensemble evaluations in the paper consider 100 instances at each dimension. Compute time t specifies the walltime for one instance of the SK problem in minutes.

n	10	12	14	16	18	20
$ E $ [MiB]	0.016	0.04	0.136	0.52	2.06	8.20
$ S $ [GiB]	0.026	0.066	0.223	0.853	3.37	13.4
t (32 cores)	0.08	0.15	0.62	2.01	12.3	74.7
t (64 cores)	0.12	0.20	0.68	1.50	8.20	54.9
t (96 cores)	0.24	0.35	4.26	1.75	7.10	28.0

2.1.2 Tool Use. Hamiltonians respectively problems are specified as text files with one term per line in format `coeff op1 i [op2 j]`. For instance:

```
-1.0 X 0
-0.7 Z 3
-1.2 Z 0 Z 1
```

Each term represents a Pauli operator acting on one or two qubits (while it is possible to consider X and Y interactions, QUBO-based formulations and standard Ising annealing problems will only resort to using Z operators, and use X operators only implicitly for the initial Hamiltonian). The format is deliberately simple, and designed to easily interface with all available mechanisms for generating QUBO problems, regardless of previously utilised programming language or quantum framework. The most relevant parameters of the tool are collected in Tab. 2.

2.1.3 Derived Quantities and Post-Processing. Primarily, the simulation produces eigenvalues $E_n(s)$ for $n = 0, \dots, k-1$, and the corresponding eigenvectors. For each eigenstate $|E_n\rangle$, the tool computes the expected value of each spin $\langle Z_i \rangle = \langle E_n(s) | Z_i | E_n(s) \rangle$ (corresponding to the local magnetisation, which is the expected value of a binary optimisation variable with values $\{-1, 1\}$), and the correlation between any pair of two spins $\langle Z_i Z_j \rangle = \langle E_n(s) | Z_i Z_j | E_n(s) \rangle$. These quantities provide insight into the structure of quantum states and their evolution on a path towards the solution for a combinatorial optimisation problem, as we discuss below.

Table 2: Main parameters of the simulation tool. Standard parameters that set output paths etc., as well as parameters that control options like tolerance settings for the employed numerical techniques are not documented here, but are available through the source code respectively the accompanying documentation.

Parameter	Description
-N	Number of qubits
-nev	Number of eigenpairs to compute
-s_start, -s_end	Range of interpolation s
-s_steps	Number of sampling points
-HI_file, -HP_file	Files for driver and problem Hamiltonian
-auto_generate_hi	Generate $H_I = -\Gamma \sum_i X_i$
-hi_gamma	Strength of driver Hamiltonian
-track_by_overlap	Enable eigenstate tracking
-track_observables	Compute $\langle Z_i \rangle$
-track_zz_correlations	Compute $\langle Z_i Z_j \rangle$
-save_eigenvectors	Store eigenvectors

Based on energy eigenvectors, a post-processing step (deliberately implemented as a second flexible step after the compute-intensive diagonalisation) included in the tool can, among others, compute energy gaps and matrix elements

$$H_{mn}(s) = \langle E_m(s) | \hat{H}(s) | E_n(s) \rangle, \quad (10)$$

$$M_{mn}(s) = \langle E_m(s) | \partial/\partial s \hat{H}(s) | E_n(s) \rangle, \quad (11)$$

where $\partial/\partial s \hat{H}(s) = \hat{H}_P - \hat{H}_I$ for the annealing Hamiltonian with linear schedule. These quantities have various uses in analysing the performance of annealing approaches. For instance, they are required when the type of phase transitions that for specific problems in an annealing process are determined, and are – in particular M_{10} – central to estimating minimally required annealing times that ascertain the desired adiabaticity. Likewise, it is possible to construct *reduced models* like *two-level models* based on (E_0, E_1, M_{10}) (so-called Landau-Zener models that are attractive for their ultimate simplicity, but usually insufficiently accurate for annealing dynamics), or more general *multi-level transport models* that use a truncated subspace to describe the effective dynamics using a simplified surrogate description (see Ref. [1, 2, 42, 48]).³ From such models, it is then possible to compute, for instance, the ground-state probability $p_{gs}(T)$ or Time-to-solution (TTS) [40]. Such metrics can connect spectral properties to algorithmic performance.

Note, however, that the purpose of our tool and this paper is *not* to construct and evaluate such models in detail – this is left for later work. Here, we set the stage for obtaining all required information for quantum problems relevant for data management.

³In the instantaneous energy eigenbasis, the evolution is known to obey:

$$i \frac{d}{dt} c_m(t) = E_m(s) c_m(t) - i s \sum_n M_{mn}(s) c_n(t). \quad (12)$$

This provides a low-dimensional approximation of the system dynamics; all required information to construct such models is available from the tool results.




Figure 2: Distribution of minimum gap values for two subject problems: Multi-Query Optimisation from data management as a combinatorial optimisation problem of domain interest, and the Sherrington-Kirkpatrick spin-glass problem as canonical archetype for hard annealing problems.

2.2 Multi-Query Optimisation

Multiple query optimisation (MQO) concerns the identification of a plan configuration that minimises the aggregate execution cost of a batch of queries evaluated jointly [60]. In several formulations, the problem is NP-hard [63]. A naïve optimiser may greedily select the individually cheapest plan for each query, thereby neglecting overlaps and potential cost-sharing opportunities across queries. Such locally optimal choices can be globally suboptimal. By contrast, MQO explicitly exploits common sub-expressions and inter-plan cost savings, seeking a configuration that minimises total cost while balancing individual plan expenses against shared savings.

Recent work by Schönberger *et al.* [59] demonstrates an approach to MQO based on a QUBO formulation, targeting quantum annealers and quantum-inspired hardware, with encouraging empirical results on the latter. As the problem can be scaled down to instances that are tractable with reasonable effort in our toolbox, and as the empirical results suggest that the problem is a good candidate for a neither too easy nor excessively hard problem in a QUBO-base formulation, we use it as reference case.

Formally, an MQO instance comprises: (i) a set of queries Q , (ii) a set of candidate execution plans P , and (iii) a set of cost-saving opportunities S . Each query $q \in Q$ is associated with a subset $P_q \subseteq P$ of mutually exclusive plans. Each plan $p_i \in P$ incurs a cost $c_i > 0$, which may be reduced through a saving $s_{p_i, p_j} \geq 0$ when combined with another plan p_j from a different query. We denote instances as MQO $|Q|/|P|, d$, where d captures the density of cost-sharing opportunities given by the fraction of cost savings featured by an MQO instance over all possible cost savings. For example, MQO 5/3, 0.25 describes five queries, three plans per query, and a sharing density of 0.25.

EXAMPLE 2.1. Consider an MQO scenario featuring four queries with plans $P_{q_1} = (p_1, p_2)$, $P_{q_2} = (p_3, p_4)$, $P_{q_3} = (p_5, p_6)$ and $P_{q_4} = (p_7, p_8)$. Further, let their execution costs be $c_1 = 9, c_2 = 10, c_3 = 9, c_4 = 10, c_5 = 11, c_6 = 9, c_7 = 14, c_8 = 9$, and the respective cost savings $s_{p_1, p_3} = 1, s_{p_1, p_4} = 1, s_{p_2, p_3} = 1, s_{p_2, p_4} = 5, s_{p_2, p_7} = 5, s_{p_4, p_5} = 5, s_{p_5, p_7} = 5, s_{p_5, p_8} = 1, s_{p_6, p_7} = 1, s_{p_6, p_8} = 1$.

Greedily, a query optimiser may choose the locally cheapest plan for each query, resulting in the overall plan selection $P_{gr} = (p_1, p_3, p_6, p_8)$. Accounting for costs savings $s_{p_1, p_3} = 1$ and $s_{p_6, p_8} = 1$, we obtain total costs $C(P_{gr}) = 9 + 9 + 9 + 9 - 1 - 1 = 34$. However, by directly accounting for such cost saving opportunities during the search process, the optimiser instead arrives at the globally optimal plan configuration, given by $P_{opt} = (p_2, p_4, p_5, p_7)$, yielding total execution cost $C(P_{opt}) = 10 + 10 + 11 + 14 - 5 - 5 - 5 - 5 = 25$.

As our goal in this paper is to show *approaches* how to set data management problems in perspective with the performance of quantum annealing, we deliberately leave an actual analysis in more depth, as well as the consideration of additional DBMS problems, to follow-up work for which this paper sets the stage.

3 Methods & Approach

3.1 Numerical Analysis

The Hilbert space of an n -qubit system has dimension 2^n , which grows exponentially with increasing problem dimension. Explicit matrix representations of $H(s)$ are therefore infeasible beyond very modest system sizes; even the largest supercomputers that are currently available can only simulate quantum systems with up to 50 qubits [47]. Instead, the computational problem is formulated as:

- Given H_I and H_P in operator form, compute the lowest k eigenvalues and eigenvectors of $H(s)$ for a set of values s_1, \dots, s_M (typically, we choose $k \in [6, 8]$ to balance accuracy with computational economy).
- Extract derived quantities such as instantaneous energy gaps $\Delta_{10}(s) = E_1(s) - E_0(s)$, expectation values $\langle Z_i \rangle$ and correlations $\langle Z_i Z_j \rangle$, and matrix elements $\langle E_m(s) | \hat{H} | E_n(s) \rangle$ and $\langle E_m(s) | \partial_s \hat{H} | E_n(s) \rangle$. All these are relevant to determine qualities and behaviour of the annealing process for a given combinatorial optimisation problem.

This leads to a sequence of large-scale sparse eigenvalue problems, that must be solved as efficiently as possible. In particular, the Hamiltonian is never materialised explicitly as a dense matrix. Instead, it is represented as a *matrix-free operator* implemented using PETSc's MatShell abstraction. The action of $\hat{H}(s)$ on a state vector is computed on-the-fly by applying Pauli operators to bit-string representations of basis states. This avoids the $O(2^{2N})$ memory requirement of dense matrices and reduces storage to $O(N)$ per term. Our implementation is mainly based on the library PETSc [4, 5] for distributed vectors, parallel communication, and matrix abstractions, and SLEPc [25–27, 49] for iterative eigenvalue solvers. Eigenpairs are computed using the Krylov-Schur subspace method [61], which is well-suited for extracting a small number of extremal eigenvalues [21]. In our simulation examples, we computed the first eight eigenvalues (in the sense of smallest algebraic value), together with the associated eigenvectors. To ensure continuity of eigenstates across s , an *overlap-based tracking* scheme is employed: We match eigenvectors obtained at discrete time s_k

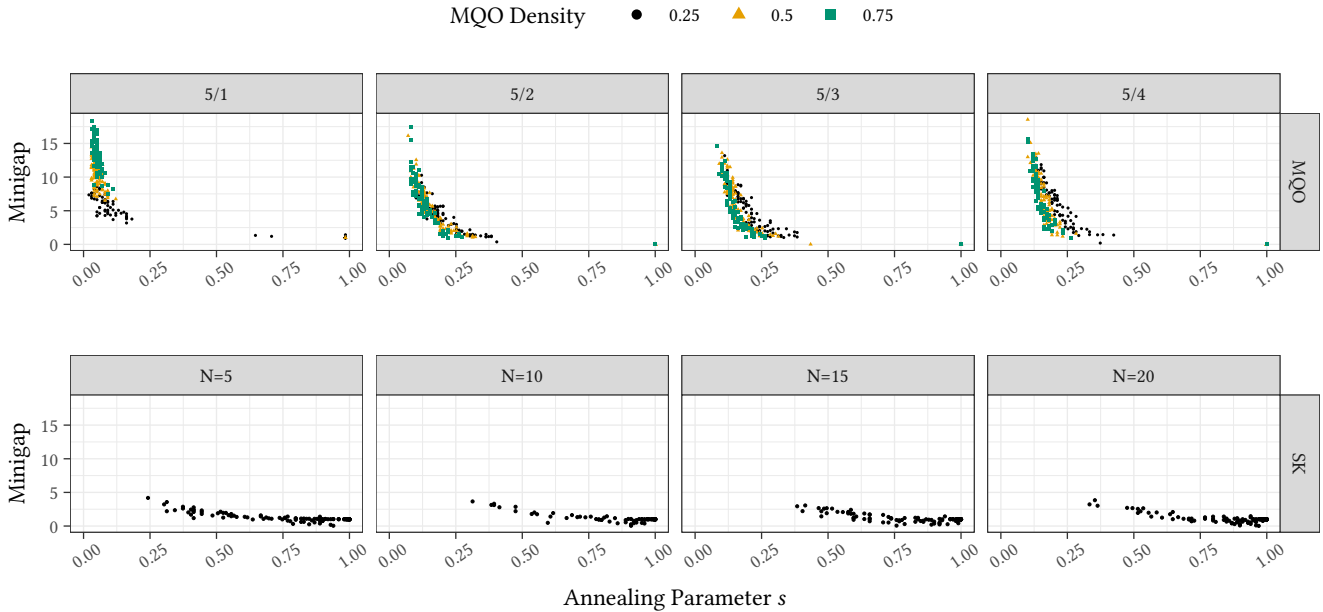


Figure 3: Minimum gap size versus temporal occurrence in the annealing process (at interpolation factor s) for 100 random instances each of SK+RF and MQO. Labels N/M for MQO denotes N queries and M plans per query; qubit counts are vertically identical across panels.

to those at discrete time s_{k-1} by maximising overlaps. This avoids artificial discontinuities from crossings or solver reordering.

3.2 Reference Problems

Apart from the key subject problem of multi-query optimisation that we use to illustrate tool capabilities, our goal is to relate hardness and complexity of quantum problem formulations in data management to existing results. Therefore, we have chosen a few reference problems with known properties that have been investigated in the literature for sometimes over decades. All problems are formulated as instances of an Ising Hamiltonian $\hat{H} = \sum_{i<j} J_{ij}s_i s_j + \sum_i h_i s_i$ with $s_i \in \{1, +1\}$ that corresponds directly to the cost function minimised by a quantum annealer at the end of the annealing schedule. Different problems are encoded by suitable choosing J_{ij} and h_i ; a direct relation between pseudo-Boolean functions [54, 55] and Ising Hamiltonians is obvious.

Ferromagnetic Ising Model. The ferromagnetic Ising model (note that different conventions for terminology are in use in the literature) serves as a tractable reference with well-understood scaling properties. We consider N spins governed by the Hamiltonian

$$H_{\text{FIM}} = -J \sum_{(i,j) \in E} s_i s_j - h \sum_i s_i. \quad (13)$$

All couplings $J > 0$ are identical and favour alignment of spins. Local terms h are identical for every spin.

The ground states without local fields (*i.e.*, $h = 0$) are the two fully aligned configurations $\mathbf{s} = (+1, \dots, +1)$ and $\mathbf{s} = (-1, \dots, -1)$, with energy $E_0 = -J|E|$. A small local bias h_i is used to break the

symmetry between these two, which also eliminates some computational challenges. Excitations correspond to domain walls, and their energy cost scales with the boundary size of flipped regions. For a complete graph, the energy gap between the ground state and first excited state scales as $\Theta(N)$, yielding a simple energy landscape without local minima.

From a quantum annealing perspective, the problem exhibits a second-order phase transition with a polynomially closing minimum gap, making this problem computationally easy. Consequently, it serves as a sanity check for correct scaling behaviour and for validating that the simulation reproduces known analytical results.

Sherrington-Kirkpatrick type models. The Sherrington-Kirkpatrick (SK) model with random external field (SK+RF; we only refer to this scenario in the following) is a paradigmatic problem. The Hamiltonian is defined as:

$$H_{\text{SK}} = - \sum_{i<j} J_{ij} s_i s_j - \sum_k h_k s_k, \quad (14)$$

where the couplings J_{ij} and the local fields h_k are drawn from a Gaussian or uniform distribution.

Appropriate normalisation ensures an extensive energy (*i.e.*, $E = \Theta(N)$). Unlike the ferromagnetic case, the interactions are frustrated: for a typical triple (i, j, k) , it is impossible to satisfy all pairwise interactions simultaneously. This leads to a highly rugged energy landscape.

The SK model exhibits well-studied physical phenomena like a phase transition at a critical, separating a paramagnetic phase from a glassy phase characterised by replica symmetry breaking. From

an optimisation viewpoint, this corresponds to a proliferation of local minima separated by extensive barriers.

In the quantum setting, the problem is believed to exhibit exponentially small minimum spectral gaps in the worst case, leading to exponentially long annealing times. The combination of disorder, frustration, and complex phase structure results in an archetypical hard benchmark problem for classical and quantum optimisation.

Hamming Weight Problem. As a complementary control problem, we consider a Hamming-weight minimisation instance, which has a trivial structure and a constant spectral gap under quantum annealing. The Hamiltonian $H_{\text{HW}} = \sum_{i=1}^N \frac{1-s_i}{2}$ is purely local; the ground state is the all-+1 configuration, with energy 0, and each spin flip increases the energy by exactly one unit. The spectrum is fully determined by the Hamming weight of the configuration: $E = |\{i \mid s_i = -1\}|$. For quantum annealing, the system decomposes into N independent single-qubit problems. The minimum spectral gap is constant and does not shrink with N , implying that the annealing time required to reach the ground state with high probability is independent of problem size. This problem provides a correctness and calibration benchmark; it is included in the tool, but we do not discuss the expected trivial results in this paper.




Figure 4: Energy curves obtained by point-wise sorting of s -local energy levels. These data are employed in computing quantities like the spectral gap $\Delta_{10}(s)$ (vertical gray bar), and are used in runtime analyses. For the subjecty problems, a clearer separation between (non-degenerate) energy levels for MOQ in contrast to SK is visible.

4 Results and Discussion

In discussing typical results, visualisations and interpretations obtainable from our tool, we highlight salient differences between

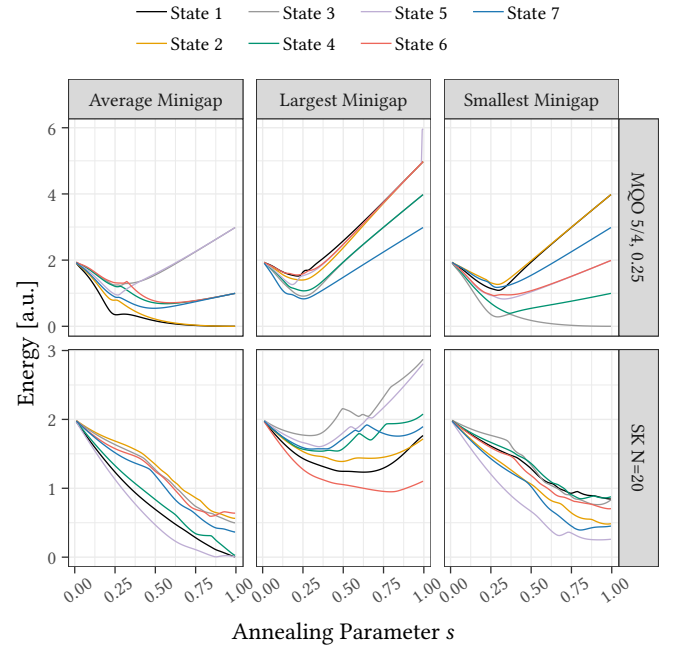


Figure 5: Energy curves for different states of the system obtained by tracking propagation from initial minimum energy overlap states. The extra effort beyond inferring point-wise sorted quantities in Fig. 4 can provide additional insights into the physical properties of an annealing process.

properties and behaviour of MOQ and SK as a known-hard reference problem, we focus on a general understanding of the meaning of the graphs, and do at this stage not proceed with a more fine-grained and detailed analysis of a specific problem relevant in data management contexts.

4.1 Instance Selection and Spectral Gaps

The properties of a given problem depend on the properties of individual instances. In general, two points of view are possible: Considering individual instances, or considering averaged ensemble properties. Representative and meaningful results require choose a representative ensemble, and to construct states with relevant properties. To ascertain the former, we use a randomly (but reproducibly) generated sample of 100 instances for each size of a given problem, and compute/visualise summaries, as in Fig. 3. Our choice of individual instances is based on the previously minimum spectral gap $\min_s \Delta_{10}(s)$ of an instance, which is a first-order proxy for minimal required annealing time (inverse quadratic) and thus also instance hardness. From the ensemble of all instances, we choose the instance with the smallest and largest minimum gap, representing a hard and easy instance, as well one ‘typical’ instance with average minimum gap in between the two boundary cases (see, e.g., Fig. 6). While this does not entail a strict correspondence with complexity-theoretical hardness, it provides an unbiased selection of relevant instances.

Fig. 2 provide an initial overview about these minimal spectra gaps across the subject instances; Fig. 3 displays the same data in a more interesting way, as it compares the minimum spectral gaps and their locations along the annealing parameter s for SK and MQO. For the Sherrington–Kirkpatrick model (bottom row), increasing the system size leads to a clear trend toward smaller minimum gaps and a shift of the dominant avoided crossings toward later stages of the anneal, consistent with the emergence of many nearly degenerate classical configurations that induce small avoided level crossings (see also Fig. 4 and the following discussion). In contrast, the multi-query optimisation instances (top row) exhibit systematically larger gaps and earlier bottlenecks, indicating that the hardest spectral rearrangements occur while the initial driver Hamiltonian \hat{H}_I still plays a significant role. Importantly, annealing dynamics are governed not just by gap size but by the structure and location of avoided crossings. These are highly instance-dependent: SK realises a dense set of late-stage near-degeneracies, while MQO shows fewer and earlier crossings with larger gaps.

4.2 State Evolution

The evolution of sorted and tracked energy branches, as defined above, is show in Fig. 4 and Fig. 5 for SK and MQO. The MQO instance exhibits a relatively smooth spectral evolution, with finite gaps throughout the annealing schedule and only mild avoided crossings, suggesting the absence of severe adiabatic bottlenecks at this system size. In contrast, the SK model displays multiple narrow avoided crossings and pronounced gap suppression near intermediate values of s , consistent with its frustrated energy landscape and dense spectrum of competing states. While these observations align with expectations that SK-type problems become exponentially hard in the large- N limit (which is the reason why we have chosen them as a reference problem), the present finite-size results primarily indicate qualitative differences in spectral structure rather than definitive asymptotic scaling. The tracked eigenstate trajectories further reveal the locations of dynamical bottlenecks that are not fully captured by point-wise sorted gaps.

The energy spectrum has important practical consequences: Knowing regions with small and large gaps (especially since different problem instances often show similar structural properties [33]) allows us to (a) construct *concrete* annealing schedules $s(t)$ that proceed faster for large and slower for small gaps. This minimises unwanted transitions (“leaks”) into higher energy levels; and (b) predict *ideal* performance on future noiseless machines as reference against existing approaches. It is (c) also the basis for HW-SW-co-design [51] approaches (e.g., adaptive error correction), or (d) semi-classical approaches based on effective physical properties.

While the spectrum is hard to compute itself, it is known to behave similarly across instances of a given problem, and knowledge for one instance can inform the choice of settings for others.

4.3 Annealing Dynamics

Fig. 6 analyses the annealing dynamics of some of our representative Hamiltonians through the quantity (usually referred to as adiabatic condition ratio or adiabatic parameter [16, 29, 48])

$$R(s) = \frac{|\langle E_1 | \partial_s \hat{H} | E_0 \rangle|}{\Delta_{10}(s)^2}, \quad (15)$$

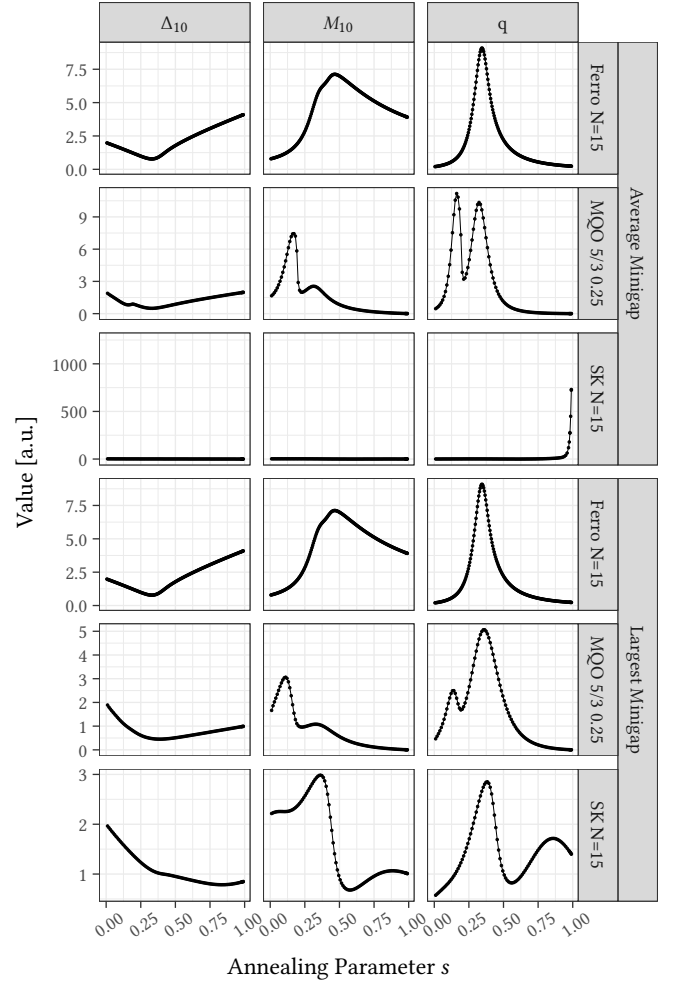


Figure 6: Temporal evolution of spectral minimum gap between ground state and first excited state energy, transition matrix element $M_{10} = |\langle E_1(s) | \partial_s \hat{H} | E_0(s) \rangle|$, and the so-called adiabatic condition ratio $q = M_{10}/\Delta_{10}^2$. The quantities provide crucial insights into dynamical and hardness properties of the annealing process; as discussed in the text.

that directly governs the local rate of diabatic transitions during quantum annealing (the plot also separately shows the constituents of the fraction; note that since the smallest minigap for SK leads to a near-singular value of R that makes it hard to reconcile the visualisation with the two other cases, we omit this instance in the plot). While the minimum spectral gap $\min_s \Delta_{10}(s)$ is frequently used as simple proxy for problem hardness, $R(s)$ provides a more informative diagnostic. For the ferromagnetic Hamiltonian, $R(s)$ exhibits a single, narrow peak roughly centred near the quantum critical point, with a modest tail extending toward larger s) (note that since we only consider a single parametrisation for a ferromagnet, the same profile is shown in both comparisons, and contains identical information). This reflects a second-order quantum phase transition accompanied by a polynomially closing gap; recall that

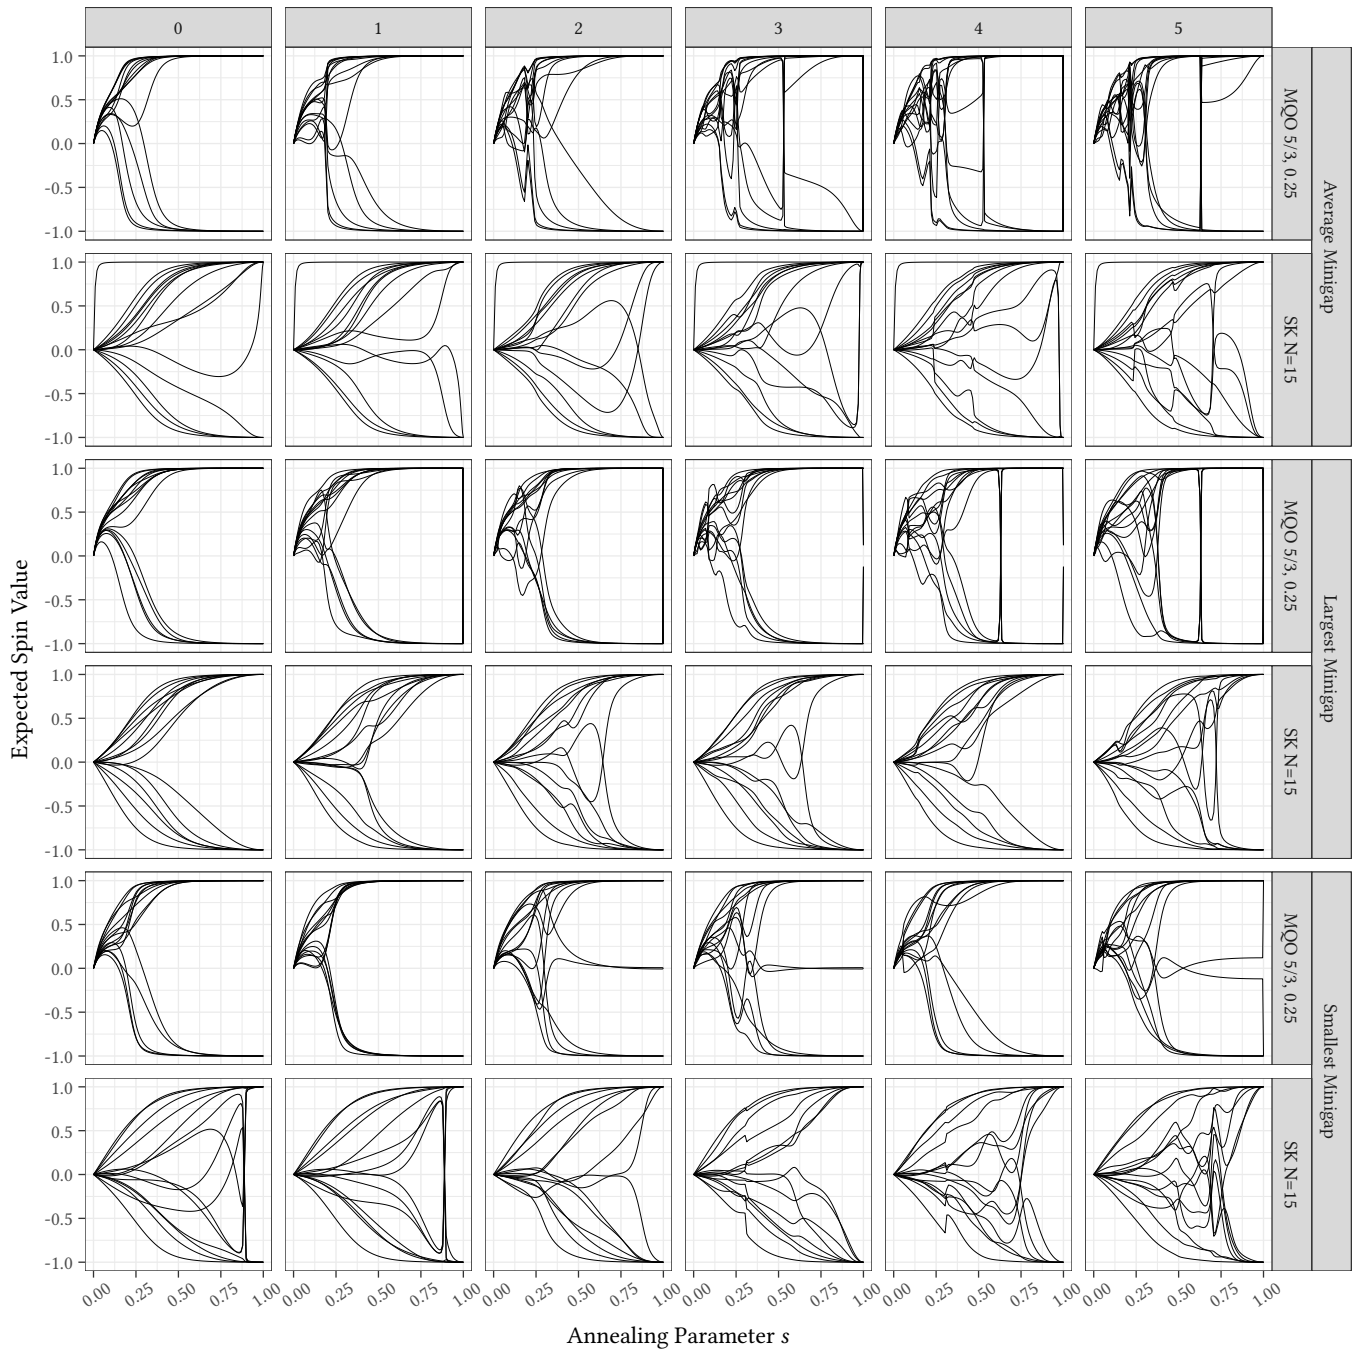


Figure 7: Spin-resolution view of temporal annealing dynamics for various instances of Multi-Query Optimisation (MQO) and the Sherrington-Kirkpatrick (SK) Hamiltonian. Each line corresponds to the expected value $\langle \sigma_i^z(s) \rangle$, the instantaneous z -magnetization of one spin. Values near 0 indicate ‘undecided qubits’, (in physical terminology, this corresponds to *frustration*), while values near ± 1 indicate qubits that have effectively converged to a classical assignment observed in the measurement process at the end of the annealing run. Smooth trajectories correspond to gradual bit fixation; abrupt collective changes signal re-organisation between competing low-energy configurations, typically near small spectral gaps. Comparing panels shows that smaller minimum gaps correlate with later, sharper, and more non-monotone spin commitments.

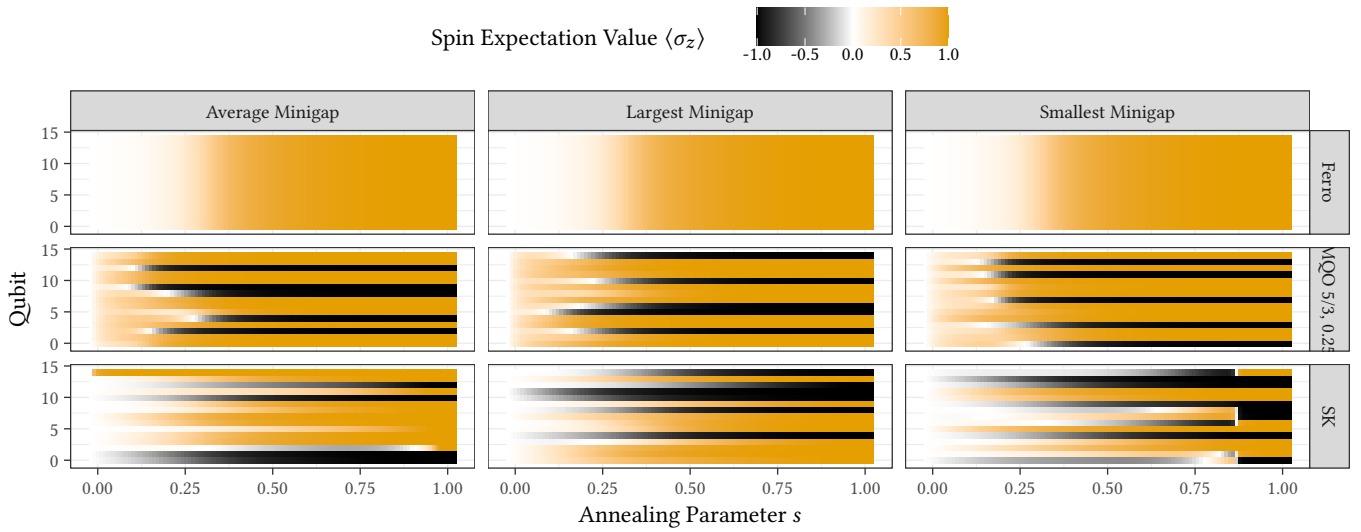


Figure 8: Development of expected values of spins settings. Each row corresponds to the expected value $\langle \sigma_i^z(s) \rangle$. Note that since the ferromagnetic Hamiltonian does not admit different instances, the data are identical across all three occurrences, which we keep as a comparison reference for an easy to solve problem on a quantum annealer.

this gap behaviour can also be analytically ascertained. The localisation of the peak implies that diabatic transitions are largely confined to a small region of the annealing schedule. This indicates that schedule optimisation on practical devices – such as locally slowing the evolution near the critical point – is effective, rendering these problem instances comparatively easy for quantum annealing. In contrast, the Sherrington–Kirkpatrick (SK) spin glass displays qualitatively richer behavior. $R(s)$ profiles remain broad and show significant weight both early and late in the anneal. This structure arises from multiple avoided crossings generated by a rugged energy landscape with many competing low-energy configurations. Diabatic transitions are not concentrated at a single bottleneck, but distributed across a wide interval of s , making the annealing process more challenging. In particular, the late-anneal features reflect residual avoided crossings among the nearly degenerate classical configurations that proliferate in the spin-glass phase.

The MQO problems display behaviour between the two extremes of a known easy and a known hard problem: $R(s)$ is generally broader and less symmetric than the ferromagnetic case, and also shows more similarity between easy and typical instances (a more detailed analysis of the consequences for data management is left to follow-up work). Nonetheless, the figure underlines that annealing hardness cannot be reduced to the minimum gap alone, and emerges from the interplay between spectral gaps and associated transition matrix elements distributed across the entire annealing path: this information is not available from straightforward empirical analysis of experimental runs on prototype devices.

4.4 Spin Dynamics and Correlations

Fig. 7 shows temporal dynamics of spins how (and thus: optimisation variables) evolve into their final measured states. As they can be in a superposition state, we use *expected* values $\langle \sigma_i^z \rangle$ over normalised annealing time. MQO instances exhibit early and largely

monotonic polarisation (*i.e.*, variables assume distinct values) of individual spins, indicating that the instantaneous ground state rapidly approaches a product state in the computational basis that gives unique measurement results. In contrast, SK instances show extended regions where $\langle \sigma_i^z \rangle \approx 0$, followed by abrupt and collective transitions. This reflects the highly frustrated, dense interaction structure of SK, where low-energy configurations differ globally and spins cannot be fixed independently. MQO behaves like a weakly coupled problem amenable to incremental decision-making, whereas SK requires coordinated global arrangements, consistent with its greater computational hardness (as usual, a detailed discussion of the consequences is left to follow-up work, as the aim of this paper is to introduce methods and computational techniques for the required analysis).

Note that we allow for deliberately tailoring the Sherrington–Kirkpatrick Hamiltonian such that one spin is effectively pinned to the up state by setting the field strength respectively coefficient appropriately. This removes exact pairwise degeneracy between globally flipped configurations. Without fixing a spin, observables odd under spin flip satisfy $\langle \sigma_i^z \rangle = 0$ in any symmetry-preserving ground state (or a thermal mixture), even when the system is effectively choosing between two opposite classical solutions. Fixing one spin selects a symmetry sector, so the remaining $\langle \sigma_i^z \rangle$ can become non-zero and reveal the structure of the chosen solution branch. Also, it can improve numerical stability and interpretability [43], as near-degenerate symmetric states can mix arbitrarily in finite-precision simulations, causing sign flips or basis-dependent behaviour from run to run. Importantly, it also adds a plausibility check to ascertain the correctness of our simulation: One spin in the average case instance simulation for SK (we apply the tailoring for this instance) immediately converges to polarisation state +1, and remains in this state throughout the complete temporal evolution because of this pinning. This is visible in the top left trajectory.

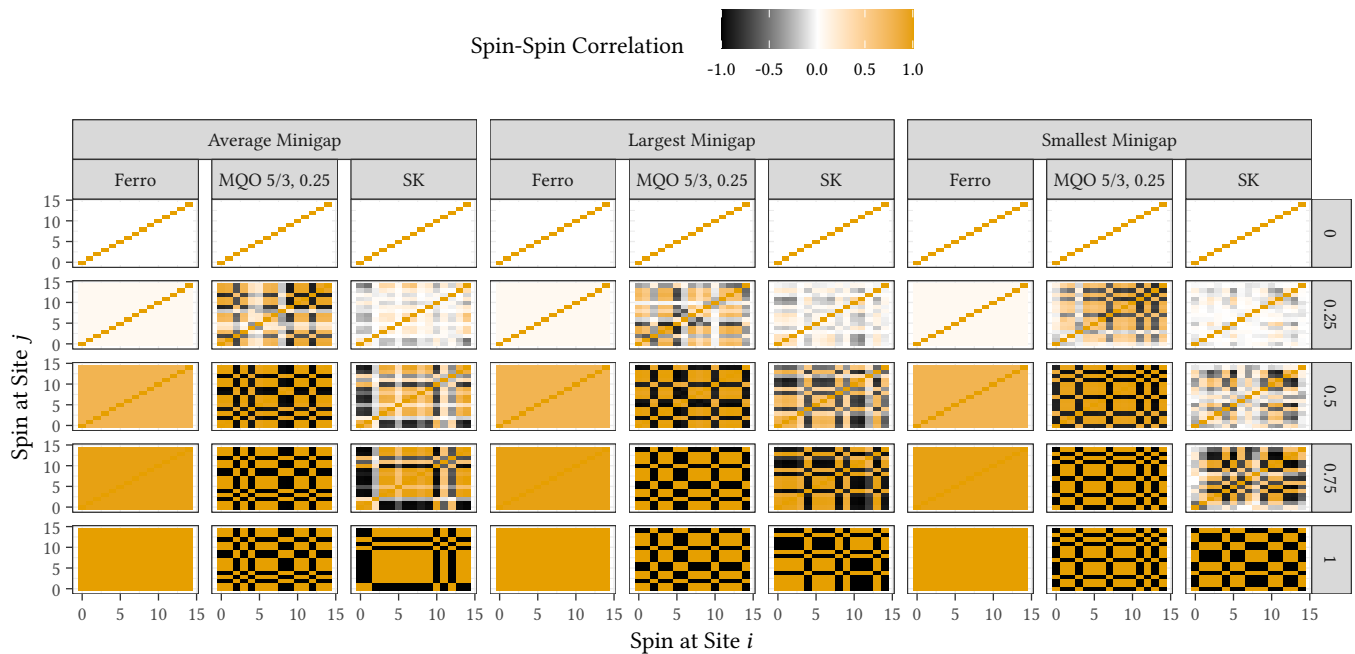


Figure 9: Spin-Spin Correlations. Each field in the grid represents the expected value $\langle \sigma_i^z \sigma_j^z \rangle$. While detail are discussed in the text, observe the particularly salient plateau with ‘undecided’ (i.e., highly fluctuating) spins that appears in white colour, corresponding to expectation value 0, for the *smallest* minimum gap instance of MQO; the corresponding phenomenon is already visible for the SK problem with an *average* minimum gap. The ferromagnet discussion in Fig. 8 applies.

Similar insights can be obtained using a different representation shown in Fig. 8 and Fig. 9: These might be more appealing to a computer science audience, as they resemble correlation and mean-value plots that are commonly used in data science. It is interesting to observe that the correlation matrices develop block/checkerboard patterns that sharpen with increasing s ; the visualisation again hints that MQO evolution is more structured than SK, but less trivial than for the ferromagnet. The blocks likely correspond to logical clusters or constraints in the optimisation problem, whereas correlations may reflect shared variables or mutual exclusion constraints.

5 Related Work

Research at the intersection of quantum computing and data management uses quantum or quantum-inspired hardware to accelerate classical database tasks such as query optimisation, transaction scheduling, schema matching, or index tuning. Cost-based query optimisation has become one of the main testbeds for quantum methods because many of its core subproblems are NP-hard and admit compact combinatorial encodings [14, 15, 63].

Join ordering is arguably the best-developed subarea within quantum query optimisation; following initial QUBO formulations [56], quantum-inspired digital annealing [57] and broader approaches from left-deep to bushy plans [45, 58, 62] have been considered, as well as quantum machine learning based methods [17, 36, 65]. Additionally, transaction scheduling [7, 8, 23], schema matching [18], index tuning and advising [24, 31], and multi-query optimisation [59, 63], have been considered; see also the survey in Ref. [11].

While latest experimental results show clear signs of beyond-classical power of quantum annealing [32], a comprehensive theory of the exact relationship between quantum phase transitions – known to be key in quantum annealing [2, 3, 66] – and parameter regimes/conditions under which quantum annealing fails if not yet known [64]. Performance of quantum annealing itself has been subject to considerable scrutiny [19, 34, 40, 41, 53], and efforts to determine the salient properties like minimum spectral gap of problems without solving the problem itself have been made [9].

6 Conclusion

We presented a computational framework for analysing combinatorial optimisation problems addressed via quantum annealing. By emphasising principled methods for evaluation, visualisation, and interpretation, we aim to provide a foundation for more rigorous and meaningful analyses of quantum approaches within the DBMS community. We anticipate that such a perspective will support more reliable assessment of potential advantages and inform future developments at the intersection of data management and quantum computing.

Acknowledgements This work was partly supported by the German Research Foundation, grant MA 9739/1-1, and by the High-Tech Agenda of the Free State of Bavaria. We also acknowledge partial support by the European Union (Project Reference 101083427) and the European Funds for Regional Development (EFRE) (Project Reference 20-3092.10-THD-105).

References

- [1] Tameem Albash and Daniel A. Lidar. 2018. Adiabatic quantum computation. *Reviews of Modern Physics* 90, 1 (Jan. 2018). doi:10.1103/revmodphys.90.015002
- [2] M. H. S. Amin. 2009. Consistency of the Adiabatic Theorem. *Physical Review Letters* 102, 22 (June 2009). doi:10.1103/physrevlett.102.220401
- [3] M. H. S. Amin and V. Choi. 2009. First-order quantum phase transition in adiabatic quantum computation. *Physical Review A* 80, 6 (Dec. 2009). doi:10.1103/physreva.80.062326
- [4] Satish Balay, Shrirang Abhyankar, Mark F. Adams, Steven Benson, Jed Brown, Peter Brune, Kris Buschelman, Emil Constantinescu, Lisandro Dalcin, Alp Dener, Victor Eijkhout, Jacob Faibussowitsch, William D. Gropp, Václav Hapla, Tobin Isaac, Pierre Jolivet, Dmitry Karpeev, Dinesh Kaushik, Matthew G. Knepley, Fande Kong, Scott Kruger, Dave A. May, Lois Curfman McInnes, Richard Tran Mills, Lawrence Mitchell, Todd Munson, Jose E. Roman, Karl Rupp, Patrick Sanan, Jason Sarich, Barry F. Smith, Hansol Suh, Stefano Zampini, Hong Zhang, Hong Zhang, and Junchao Zhang. 2025. *PETSc/TAO Users Manual*. Technical Report ANL-21/39 - Revision 3.24. Argonne National Laboratory. doi:10.2172/2998643
- [5] Satish Balay, William D. Gropp, Lois Curfman McInnes, and Barry F. Smith. 1997. Efficient Management of Parallelism in Object Oriented Numerical Software Libraries. In *Modern Software Tools in Scientific Computing*, E. Arge, A. M. Bruaset, and H. P. Langtangen (Eds.). Birkhäuser Press, 163–202.
- [6] Andreas Bayerstadler, Guillaume Becquin, Julia Binder, Thierry Botter, Hans Ehm, Thomas Ehmer, Marvin Erdmann, Norbert Gaus, Philipp Harbach, Maximilian Hess, Johannes Klepsch, Martin Leib, Sebastian Lubser, Andre Luckow, Maximilian Mansky, Wolfgang Mauerer, Florian Neukart, Christoph Niedermeier, Lilly Palackal, Ruben Pfeiffer, Carsten Polenz, Johanna Sepulveda, Tammo Sievers, Brian Standen, Michael Streif, Thomas Strohm, Clemens Utschig-Utschig, Daniel Volz, Horst Weiss, and Fabian Winter. 2021. Industry Quantum Computing Applications. *EPJ Quantum Technology* 8, 1 (11 2021). doi:10.1140/epjqt/s40507-021-00114-x
- [7] Tim Bittner and Sven Groppe. 2020. Avoiding Blocking by Scheduling Transactions Using Quantum Annealing. In *Proceedings of the 24th International Database Engineering & Applications Symposium*. 21:1–21:10. doi:10.1145/3410566.3410593
- [8] Tim Bittner and Sven Groppe. 2020. Hardware Accelerating the Optimization of Transaction Schedules via Quantum Annealing by Avoiding Blocking. *Open Journal of Cloud Computing* 7, 1 (2020), 1–21.
- [9] Tim Bode and Frank K. Wilhelm. 2024. Adiabatic bottlenecks in quantum annealing and nonequilibrium dynamics of paramagnons. *Physical Review A* 110, 1 (July 2024). doi:10.1103/physreva.110.012611
- [10] Cecilia Carbonelli, Michael Felderer, Matthias Jung, Elisabeth Lobe, Malte Lochau, Sebastian Lubser, Wolfgang Mauerer, Rudolf Ramler, Ina Schäfer, and Christoph Schroth. 2024. Challenges for Quantum Software Engineering: An Industrial Use Case Perspective. In *Quantum Software: Aspects of Theory and System Design*. Springer-Nature. doi:10.1007/978-3-031-64136-7_12
- [11] Umut Çalkıylmaz, 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. *Proceedings of the VLDB Endowment* 16, 9 (2023), 2344–2353. doi:10.14778/3598581.3598603
- [12] E. J. Crosson and D. A. Lidar. 2021. Prospects for quantum enhancement with diabatic quantum annealing. *Nature Reviews Physics* 3, 7 (May 2021), 466–489. doi:10.1038/s42254-021-00313-6
- [13] Jens Eisert and John Preskill. 2025. Mind the gaps: The fraught road to quantum advantage. arXiv:2510.19928 [quant-ph] <https://arxiv.org/abs/2510.19928>
- [14] Tobias Fankhauser, Marc E. Solèr, Rudolf M. Fuchsln, and Kurt Stockinger. 2021. Multiple Query Optimization Using a Hybrid Approach of Classical and Quantum Computing. *CoRR* abs/2107.10508 (2021). doi:10.48550/arXiv.2107.10508
- [15] Tobias Fankhauser, Marc E. Solèr, Rudolf Marcel Fuchsln, and Kurt Stockinger. 2023. Multiple Query Optimization Using a Gate-Based Quantum Computer. *IEEE Access* 11 (2023), 114031–114043. doi:10.1109/ACCESS.2023.3324253
- [16] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. 2000. Quantum Computation by Adiabatic Evolution. *arXiv preprint* (2000). arXiv:quant-ph/0001106
- [17] Maja Franz, Tobias Winker, Sven Groppe, and Wolfgang Mauerer. 2024. Hype or Heuristic? Quantum Reinforcement Learning for Join Order Optimisation. In *IEEE International Conference on Quantum Computing and Engineering (QCE 2024)*. 409–420. doi:10.1109/QCE60285.2024.00055
- [18] Kristin Fritsch and Stefanie Scherzinger. 2023. Solving Hard Variants of Database Schema Matching on Quantum Computers. *Proceedings of the VLDB Endowment* 16, 12 (2023), 3990–3993. doi:10.14778/3611540.3611603
- [19] Thomas Gabor, Sebastian Zielinski, Sebastian Feld, Christoph Roch, Christian Seidel, Florian Neukart, Isabella Galter, Wolfgang Mauerer, and Claudia Linnhoff-Popien. 2019. Assessing Solution Quality of 3SAT on a Quantum Annealing Platform. In *Quantum Technology and Optimization Problems*. Springer International Publishing, 23–35. doi:10.1007/978-3-030-14082-3_3
- [20] Martin Gogeißl, Hila Safi, and Wolfgang Mauerer. 2024. Quantum Data Encoding Patterns and their Consequences. In *Proceedings of the Workshop on Quantum Computing and Quantum-Inspired Technology for Data-Intensive Systems and Applications (Q-Data '24)*. doi:10.1145/3665225.3665446
- [21] Gene H. Golub and Charles F. Van Loan. 2013. *Matrix Computations - 4th Edition*. Johns Hopkins University Press, Philadelphia, PA. arXiv:<https://epubs.siam.org/doi/pdf/10.1137/1.9781421407944> doi:10.1137/1.9781421407944
- [22] Felix Greiwe, Tom Krüger, and Wolfgang Mauerer. 2023. Effects of Imperfections on Quantum Algorithms: A Software Engineering Perspective. In *2023 IEEE International Conference on Quantum Software (QSW)*. 31–42. doi:10.1109/QSW59989.2023.00014
- [23] 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*. 149–156. doi:10.1145/3472163.3472164
- [24] Le Gruenwald, Tobias Winker, Umut Çalkıylmaz, 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 (VLDBW 2023) - International Workshop on Quantum Data Science and Management (QDSM'23)*. <https://ceur-ws.org/Vol-3462/QDSM5.pdf>
- [25] V. Hernández, J. E. Román, and A. Tomás. 2007. Parallel Arnoldi eigensolvers with enhanced scalability via global communications rearrangement. *Parallel Comput.* 33, 7-8 (2007), 521–540.
- [26] V. Hernandez, J. E. Roman, and V. Vidal. 2003. SLEPc: Scalable Library for Eigenvalue Problem Computations. *Lect. Notes Comput. Sci.* 2565 (2003), 377–391.
- [27] Vicente Hernandez, Jose E. Roman, and Vicente Vidal. 2005. SLEPc: A scalable and flexible toolkit for the solution of eigenvalue problems. *ACM Trans. Math. Software* 31, 3 (2005), 351–362.
- [28] Torsten Hoefler, Thomas Häner, and Matthias Troyer. 2023. Disentangling Hype from Practicality: On Realistically Achieving Quantum Advantage. *Commun. ACM* 66, 5 (April 2023), 82–87. doi:10.1145/3571725
- [29] Sabine Jansen, Mary-Beth Ruskai, and Ruedi Seiler. 2007. Bounds for the adiabatic approximation with applications to quantum computation. *J. Math. Phys.* 48, 10 (Oct. 2007). doi:10.1063/1.2798382
- [30] Tadashi Kadowaki and Hidetoshi Nishimori. 2022. Greedy parameter optimization for diabatic quantum annealing. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 381, 2241 (Dec. 2022). doi:10.1098/rsta.2021.0416
- [31] Manish Kesarwani and Jayant R. Haritsa. 2024. Index Advisors on Quantum Platforms. *Proceedings of the VLDB Endowment* 17, 11 (2024), 3615–3628. doi:10.14778/3681954.3682025
- [32] Andrew D. King, Alberto Nocera, Marek M. Rams, Jacek Dziarmaga, Roeland Wiersema, William Bernoudy, Jack Raymond, Nitin Kaushal, Niclas Heinsdorf, Richard Harris, Kelly Boothby, Fabio Altomare, Mohsen Asad, Andrew J. Berkley, Martin Boschnak, Kevin Chern, Holly Christiani, Samantha Cibere, Jake Connor, Martin H. Dehn, Rahul Deshpande, Sara Ejtemaee, Pau Farre, Kelsey Hamer, Emile Hoskinson, Shuiyuan Huang, Mark W. Johnson, Samuel Kortas, Eric Ladizinsky, Trevor Lanting, Tony Lai, Ryan Li, Allison J. R. MacDonald, Gaelen Marsden, Catherine C. McGeoch, Reza Molavi, Travis Oh, Richard Neufeld, Mana Norouzpour, Joel Pasvolksy, Patrick Poitras, Gabriel Poulin-Lamarre, Thomas Prescott, Mauricio Reis, Chris Rich, Mohammad Samani, Benjamin Sheldan, Anatoly Smirnov, Edward Sterpka, Berta Trullas Clavera, Nicholas Tsai, Mark Volkmann, Alexander M. Whittaker, Jed D. Whittaker, Warren Wilkinson, Jason Yao, T. J. Yi, Anders W. Sandvik, Gonzalo Alvarez, Roger G. Melko, Juan Carrasquilla, Marcel Franz, and Mohammad H. Amin. 2025. Beyond-classical computation in quantum simulation. *Science* 388, 6743 (2025), 199–204. arXiv:<https://www.science.org/doi/pdf/10.1126/science.ado6285> doi:10.1126/science.ado6285
- [33] Tom Krüger and Wolfgang Mauerer. 2025. Out of the Loop: Structural Approximation of Optimisation Landscapes and non-Iterative Quantum Optimisation. *Quantum* 9 (Nov. 2025), 1903. doi:10.22331/q-2025-11-06-1903
- [34] Tom Krüger and Wolfgang Mauerer. 2020. Quantum Annealing-Based Software Components: An Experimental Case Study with SAT Solving. In *Proceedings of the IEEE/ACM 42nd International Conference on Software Engineering Workshops (Seoul, Republic of Korea) (ICSEW'20)*. Association for Computing Machinery, New York, NY, USA, 445–450. doi:10.1145/3387940.3391472
- [35] Gushu Li, Anbang Wu, Yunong Shi, Ali Javadi-Abhari, Yufei Ding, and Yuan Xie. 2021. On the Co-Design of Quantum Software and Hardware. In *Proceedings of the Eight Annual ACM International Conference on Nanoscale Computing and Communication (Virtual Event, Italy) (NANOCOM '21)*. Association for Computing Machinery, New York, NY, USA, Article 15, 7 pages. doi:10.1145/3477206.3477464
- [36] Hanwen Liu, Federico M. Spedalieri, and Ibrahim Sabek. 2025. A Demonstration of Q2Q: Quantum-augmented Query Optimizer. *Proceedings of the VLDB Endowment* 18, 12 (2025), 5439–5443. doi:10.14778/3750601.3750691
- [37] Jeanette Miriam Lorenz, Thomas Monz, Jens Eisert, Daniel Reitzner, Félicien Schopfer, Frédéric Barbaresco, Krzysztof Kurowski, Ward van der Schoot, Thomas Strohm, Jean Senellart, Cécile M. Perratot, Martin Knufinke, Ziyad Amodee, and

- Mattia Giardini. 2025. Systematic benchmarking of quantum computers: status and recommendations. arXiv:2503.04905 [quant-ph] <https://arxiv.org/abs/2503.04905>
- [38] Stefan Raimund Maschek, Jürgen Schwittalla, Maja Franz, and Wolfgang Mauerer. 2025. Make Some Noise! Measuring Noise Model Quality in Real-World Quantum Software. In *Proceedings of the IEEE International Conference on Quantum Software (QSW)*. arXiv:2506.03636 doi:10.1109/QSW67625.2025.00010
- [39] Wolfgang Mauerer and Stefanie Scherzinger. 2022. 1-2-3 Reproducibility for Quantum Software Experiments. In *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*. 1247–1248. doi:10.1109/SANER53432.2022.00148
- [40] Vrinda Mehta, Hans De Raedt, Kristel Michielsens, and Fengping Jin. 2025. Performance of quantum annealing for 2-satisfiability problems with multiple satisfying assignments. *Phys. Rev. A* 112 (Jul 2025), 012405. Issue 1. doi:10.1103/n7r5-s63q
- [41] Vrinda Mehta, Fengping Jin, Hans De Raedt, and Kristel Michielsens. 2021. Quantum annealing with trigger Hamiltonians: Application to 2-satisfiability and nonstoquastic problems. *Phys. Rev. A* 104 (Sep 2021), 032421. Issue 3. doi:10.1103/PhysRevA.104.032421
- [42] Eugen Merzbacher. 1997. *Quantum Mechanics, 3rd Edition*.
- [43] M. Mezard, G. Parisi, and M. Virasoro. 1986. *Spin Glass Theory and Beyond*. WORLD SCIENTIFIC. arXiv:<https://www.worldscientific.com/doi/pdf/10.1142/0271> doi:10.1142/0271
- [44] Siddharth Muthukrishnan, Tameem Albash, and Daniel A. Lidar. 2016. Tunneling and Speedup in Quantum Optimization for Permutation-Symmetric Problems. *Physical Review X* 6, 3 (July 2016). doi:10.1103/physrevx.6.031010
- [45] Nitin Nayak, Jan Rehfeld, Tobias Winker, Benjamin Warnke, Umut Çalikylmaz, 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*. 7:1–7:7. doi:10.1145/3579142.3594298
- [46] A. Pelissetto and E. Vicari. 2024. Scaling Behaviors at Quantum and Classical First-Order Transitions. In *50 Years of the Renormalization Group*. World Scientific, 437–476. doi:10.1142/9789811282386_0027
- [47] Hans De Raedt, Jiri Kraus, Andreas Hertent, Vrinda Mehta, Mathis Bode, Markus Hrywniak, Kristel Michielsens, and Thomas Lippert. 2025. Universal Quantum Simulation of 50 Qubits on Europe's First Exascale Supercomputer Harnessing Its Heterogeneous CPU-GPU Architecture. arXiv:2511.03359 [quant-ph] <https://arxiv.org/abs/2511.03359>
- [48] Jérémie Roland and Nicolas J. Cerf. 2002. Quantum search by local adiabatic evolution. *Physical Review A* 65, 4 (March 2002). doi:10.1103/physreva.65.042308
- [49] J. E. Roman, F. Alvarruiz, C. Campos, L. Dalcin, P. Jolivet, and A. Lamas Daviña. 2023. Improvements to SLEPC in releases 3.14–3.18. *ACM Trans. Math. Software* 49, 3 (2023), 29:1–29:11.
- [50] Hila Safi, Medina Bandic, Christoph Niedermeier, Carmen G. Almudever, Sebastian Feld, and Wolfgang Mauerer. 2025. Stacking the odds: full-stack quantum system design space exploration. *EPJ Quantum Technology* 12, 1 (Oct. 2025). doi:10.1140/epjqt/s40507-025-00413-7
- [51] Hila Safi, Karen Wintersperger, and Wolfgang Mauerer. 2023. Influence of HW-SW-Co-Design on Quantum Computing Scalability. In *IEEE International Conference on Quantum Software (QSW)*. 104–115. doi:10.1109/QSW59989.2023.00022
- [52] Giulia Salatino, Maximilian Matzler, Annarita Scocco, Procolo Lucignano, and Gianluca Passarelli. 2025. Noise effects on diabatic quantum annealing protocols. *Physical Review A* 112, 2 (Aug. 2025). doi:10.1103/x9hw-xhvj
- [53] Irmi Sax, Sebastian Feld, Sebastian Zielinski, Thomas Gabor, Claudia Linnhoff-Popien, and Wolfgang Mauerer. 2020. Approximate approximation on a quantum annealer. In *Proceedings of the 17th ACM International Conference on Computing Frontiers (Catania, Sicily, Italy) (CF '20)*. Association for Computing Machinery, New York, NY, USA, 108–117. doi:10.1145/3387902.3392635
- [54] Lukas Schmidbauer, Elisabeth Lobe, Ina Schaefer, and Wolfgang Mauerer. 2026. It's Quick to be Square: Fast Quadratisation for Quantum Toolchains. *ACM Transactions on Quantum Computing* (March 2026). doi:10.1145/3800943
- [55] Lukas Schmidbauer, Carlos A. Riofrio, Florian Heinrich, Vanessa Junk, Ulrich Schwenk, Thomas Husslein, and Wolfgang Mauerer. 2025. Path Matters: Industrial Data Meet Quantum Optimization. In *2025 IEEE International Conference on Quantum Computing and Engineering (QCE)*, Vol. 01. 2101–2111. doi:10.1109/QCE65121.2025.00230
- [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. doi: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. doi:10.14778/3632093.3632112
- [58] Manuel Schönberger, Immanuel Trummer, and Wolfgang Mauerer. 2023. Quantum Optimisation of General Join Trees. In *Joint Proceedings of Workshops at the 49th International Conference on Very Large Data Bases (VLDBW 2023) – International Workshop on Quantum Data Science and Management (QDSM'23)*. <https://ceur-ws.org/Vol-3462/QDSM2.pdf>
- [59] Manuel Schönberger, Immanuel Trummer, and Wolfgang Mauerer. 2025. Large-Scale Multiple Query Optimisation with Incremental Quantum(-Inspired) Annealing. *Proc. ACM Manag. Data* 3, 4, Article 253 (Sept. 2025), 25 pages. doi:10.1145/3749171
- [60] Timos K. Sellis. 1988. Multiple-Query Optimization. *ACM Trans. Database Syst.* 13, 1 (mar 1988), 23–52. doi:10.1145/42201.42203
- [61] G. W. Stewart. 2002. A Krylov–Schur Algorithm for Large Eigenproblems. *SIAM J. Matrix Anal. Appl.* 23, 3 (Jan. 2002), 601–614. doi:10.1137/s0895479800371529
- [62] 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)*. 1–2. doi:10.1145/3736393.3736690
- [63] 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. doi:10.14778/2947618.2947621
- [64] Matthias Werner, Artur García-Sáez, and Marta P. Estarellas. 2023. Bounding first-order quantum phase transitions in adiabatic quantum computing. *Phys. Rev. Res.* 5 (Dec 2023), 043236. Issue 4. doi:10.1103/PhysRevResearch.5.043236
- [65] Tobias Winker, Umut Çalikylmaz, 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*. 5:1–5:7. doi:10.1145/3579142.3594299
- [66] A. P. Young, S. Knysh, and V. N. Smelyanskiy. 2010. First-Order Phase Transition in the Quantum Adiabatic Algorithm. *Physical Review Letters* 104, 2 (Jan. 2010). doi:10.1103/physrevlett.104.020502
- [67] Tao Yue, Wolfgang Mauerer, Shaikat Ali, and Davide Taibi. 2023. Challenges and Opportunities in Quantum Software Architecture. In *Software Architecture: Research Roadmaps from the Community*. Springer Nature Switzerland, Cham, 1–23. doi:10.1007/978-3-031-36847-9_1