

Skipper: Improving the Reach and Fidelity of Quantum Annealers by Skipping Long Chains

Ramin Ayanzadeh
Georgia Institute of Technology
Atlanta, USA

Moinuddin Qureshi
Georgia Institute of Technology
Atlanta, USA

Abstract

Quantum Annealers (QAs) operate as single-instruction machines, lacking a SWAP operation to overcome limited qubit connectivity. Consequently, multiple physical qubits are *chained* to form a program qubit with higher connectivity, resulting in a drastically diminished effective QA capacity by up to 33x. We observe that in QAs: (a) chain lengths exhibit a power-law distribution, a few *dominant chains* holding substantially more qubits than others; and (b) about 25% of physical qubits remain unused, getting isolated between these chains. We propose *Skipper*, a software technique that enhances the capacity and fidelity of QAs by skipping dominant chains and substituting their program qubit with two readout results. Using a 5761-qubit QA, we demonstrate that Skipper can tackle up to 59% (Avg. 28%) larger problems when eleven chains are skipped. Additionally, Skipper can improve QA fidelity by up to 44% (Avg. 33%) when cutting five chains (32 runs). Users can specify up to eleven chain cuts in Skipper, necessitating about 2,000 distinct quantum executable runs. To mitigate this, we introduce *Skipper-G*, a greedy scheme that skips sub-problems less likely to hold the global optimum, executing a maximum of 23 quantum executables with eleven chain trims. Skipper-G can boost QA fidelity by up to 41% (Avg. 29%) when cutting five chains (11 runs).

1 Introduction

Quantum computers (QCs) have the potential to solve certain problems beyond the capabilities of classical computers [7, 45, 78, 93, 98]. Two main types of QCs exist: digital machines, exemplified by industry leaders like IBM [26], Google [32], IonQ [41], and Quantinuum [36], and analog devices such as superconducting *Quantum Annealers* (QAs) developed by D-Wave [39], as well as neutral atom platforms by QuEra [80] and PASQAL [65]. Both digital and analog QCs have polynomial equivalent computing power [2, 4]. For instance, QAs have demonstrated their potential in tackling real-world applications such as finance [29], drug discovery [57], cryptography [38, 77], Boolean Satisfiability (SAT) [14, 15, 17, 83], planning and scheduling [40, 81, 91, 92], linear algebra [64], and signal processing [16, 18], extending beyond application-specific acceleration.

While both QC types are accessed via the cloud [6, 39, 52], their operation models and design trade-offs differ significantly [10]. In digital QCs (namely the gate-based or circuit model quantum computing), as shown in Fig. 1(a), qubits

undergo a scheduled sequence of quantum operations defined by the quantum algorithm to directly manipulate their states [60]. Conversely, as shown in Fig. 1(b), analog QCs operate as single-instruction systems, where the qubit environment is incrementally modified based on the evolution of a physical system, called *Hamiltonian*, thereby allowing natural qubit evolution and indirect state alteration [4, 10, 51].




Figure 1. Operation Model of QCs: (a) Digital QCs execute compiled quantum circuits. (b) Analog QAs execute the embedded problem Hamiltonian. *Operating as single-instruction machines, analog QAs do not incorporate quantum circuits.*

Full connectivity of qubits at scale is infeasible. In digital QCs, compilers introduce SWAP operations to make physical qubits adjacent [59, 90, 103]. Conversely, analog QCs cannot apply operations to qubits, thus preventing the use of SWAPs for qubit routing. Instead, QAs employ *embedding* [21, 69, 71, 75, 102] where multiple physical qubits are *chained* (or entangled) to represent a program qubit with higher connectivity, as shown in Fig. 2(a). Compiling quantum circuits in digital QCs preserves qubit utilization (1-to-1 mapping between program and physical qubits), however, embedding in QAs can substantially increase physical qubit utilization [10]. For instance, the 5761-qubit QA can accommodate up to 177 program qubits with all-to-all connectivity, highlighting nearly 33x reduced *logical capacity*.

Given that the hardware graph remains fixed after fabrication, QAs’ logical capacity is primarily determined by the topology of the problem graph. Real-world applications typically involve irregular “Power-Law” graphs [1, 25, 30, 31, 37, 54, 66], and Barabasi–Albert (BA) graphs are widely considered representative of such real-world graphs [5, 19, 20, 34, 43, 47, 97, 101, 102]. Figure 2(b) illustrates the largest embeddable BA graphs on a 5761-qubit QA, ranging from sparse (with attachment factor $m = 1$, BA-1) to nearly fully connected (with $m = 6$, BA-6) structures. As m increases linearly, the logical capacity experiences a superpolynomial reduction, converging to the 177-node fully connected graph.




Figure 2. (a) Embedding seven program qubits (Q_i) onto a 5×7 grid of physical qubits. (b) Max embeddable Barabasi-Albert (BA) graphs on a 5761-qubit QA device for different preferential attachment factors (m), ranging from sparse BA-1 ($m = 1$) to highly dense BA-6 ($m = 6$) structures.

We observe that chain lengths in QAs follow a “Power-Law” distribution, where a few *dominant chains* are significantly longer than most other chains (see section 4.1 for more information). Moreover, we observe that a significant portion of physical qubits, nearly 25%, remain unused as they become trapped in long chains. Furthermore, we observe that long chains can reduce the fidelity of QAs too. The qubits within a chain might take different values post-measurement, called *broken chains*. Broken chains can negatively impact QAs’ reliability, and longer chains are more likely to break.

In this study, we aim to improve the capacity and fidelity of QAs through eliminating dominant chains, as they account for a substantial portion of qubit utilization and are the main reason for isolating physical qubits. We propose *Skipper*, which *prunes* these chains by removing their corresponding program qubits and replacing them with two possible measurement outcomes: -1 and $+1$. By eliminating a dominant chain, Skipper accomplishes two objectives: (a) releasing physical qubits previously used within pruned chains, and (b) releasing all qubits previously trapped with the pruned chain. This can enable us to use all released physical qubits to accommodate more program qubits.

However, identifying and pruning dominant chains is non-trivial. Chains are formed post-embedding. First, when a (long) chain is eliminated, the remaining embedding is likely not to be optimum, necessitating re-embedding the new problem to maximize the reliability of QAs. Embedding itself is nontrivial, as it can take several hours for problems at scale. Moreover, embedding techniques are heuristic, and they may fail to find an embedding successfully for a problem, requiring multiple attempts. Second, pruning the longest chain can change the position of the second-longest chain when re-embedding the problem, necessitating an embedding for every pruned chain. To this end, Skipper adopts a greedy approach to prune c chains by sorting program qubits based on their degree and removing the top c qubits simultaneously. We observe that this greedy approach exhibits desirable, near-optimal behavior for $c \geq 5$ chain cuts.

Importantly, the number of chain cuts in Skipper is user-defined; the system allows for a maximum of eleven chains to be cut, and this does not scale with the problem size, offering flexibility within the user’s budget constraints. Each chain cut bifurcates the search space of the initial problem; therefore, trimming eleven chains can lead to up to 2048 sub-problems, and Skipper examines all of them to ensure guaranteed recovery. Our experiments with a 5761-qubit QA by D-Wave demonstrate that Skipper can address up to 59% (Avg. 28.3%) larger problems when up to eleven chains are trimmed. Additionally, Skipper can significantly enhance QA fidelity by up to 44.4% (Avg. 33.1%), when trimming up to five chains and running 32 quantum executables.

Skipper is inspired by FrozenQubits [9]. Skipper enhances both the capacity and fidelity of analog QAs. However, FrozenQubits has a negligible impact on the capacity of digital QCs, where one program qubit is represented with one physical qubit. Furthermore, FrozenQubits’ performance diminishes as graph density increases, whereas Skipper effectively handles graphs ranging from sparse to dense structures.

The quantum cost of Skipper can present affordability challenges for certain users. We introduce *Skipper-G*, a greedy approach that bypasses sub-spaces less likely to include the global optimum. Consequently, it runs at most 23 quantum executables, compared to the 2048 required by Skipper for trimming up to eleven chains. It is worth noting that Skipper-G is proposed to improve QA fidelity, with its effect on increasing capacity being negligible. Our experiments demonstrate that Skipper-G can boost QA fidelity by up to 40.8% (Avg. 29.2%), with five chain cuts and 11 runs.

Overall, this paper makes the following contributions:

1. We show that in QAs, the chain length exhibits a “Power-Law” distribution, with a few dominant chains having significantly more qubits. Moreover, we demonstrate that approximately 25% of physical qubits remain unused as they become trapped within long chains.
2. We introduce *Skipper* that enhances the capacity and reliability of QAs by cutting dominant chains, thereby addressing up to 59% (Avg. 28.3%) larger problems and improving QA fidelity by up to 44.4% (Avg. 33.1%), when up to eleven and five chains are cut, respectively. To our knowledge, Skipper is the first proposal to simultaneously enhance both the capacity and fidelity of QAs.
3. We demonstrate that the quantum cost of Skipper in enhancing QA fidelity can be substantially reduced (to only 23 runs, compared to over 2000 runs in Skipper) by bypassing sub-spaces unlikely to contain optimal solutions.
4. We propose *Skipper-G*, a greedy scheme that enhances QA fidelity by up to 40.8% (Avg. 29.2%), with five chain cuts and only 11 runs (compared to 32 runs in Skipper).

2 Background and Motivation

2.1 Quantum Computers: Digital vs. Analog

QCs fall into two categories: digital and analog. Digital QCs, like those from IBM and Google, apply precise quantum operations—defined by the quantum algorithm—to qubits in order to directly manipulate their state [60]. Conversely, analog QCs, like those from D-Wave and QuEra, do not directly manipulate the state of qubits. Instead, they apply precise changes—defined by the quantum program—to the environment in which the qubits reside, allowing the qubits to evolve and change their states naturally [4, 10].

2.2 Quantum Annealers: Analog Quantum Accelerators

Quantum annealing is a meta-heuristic for tackling optimization problems that runs on classical computers. *Quantum Annealers (QAs)* are a form of analog QCs that can sample from the ground state (the configuration with the lowest energy value) of a physical system, called Hamiltonian [4, 10, 11]. QAs by D-Wave are single-instruction optimization accelerators that can only sample from the ground state of the following problem Hamiltonian (or Ising model):

$$\mathcal{H}_p := \sum_i \mathbf{h}_i z_i + \sum_{i \neq j} J_{ij} z_i z_j \quad (1)$$

acting on spin variables $z_i \in -1, +1$, where $\mathbf{h}_i \in \mathbb{R}$ and $J_{ij} \in \mathbb{R}$ are linear and quadratic coefficients, respectively [10].

2.3 Operation Model of Single-Instruction QAs

QAs operate as single-instruction computers, and during each execution trial, they only draw a single sample to approximate the global minimum of (1). Therefore, we *cast* real-world problems into Hamiltonians, where $\mathbf{h}$ and J are defined in such a way that its global minimum represents the optimal solution to the problem at hand [4, 10]. The abstract problem Hamiltonian is then *embedded* into the connectivity map of the QA hardware to generate an executable Quantum Machine Instruction (QMI) [23, 85]. Casting and embedding in QAs are akin to designing and compiling quantum circuits in digital QCs, respectively (Fig. 1). The QMI is executed for several trials, and the outcome with the lowest objective value is deemed as the ultimate result [10].

2.4 Anneal Time: Current Technological Barriers

As the energy gap between the global minimum and the adjacent higher state diminishes linearly, the required annealing time for successful adiabaticity grows exponentially [4, 27], surpassing the limits of contemporary QAs [99]. Nonetheless, QAs, akin to other QCs, are advancing; subsequent generations are expected to bypass present technological constraints. Specifically, incorporating XX terms into the time-dependent Hamiltonian can ebb the annealing time scaling from exponential to linear [61].




Figure 3. Embedding example.

2.5 Embedding for QAs

The connectivity of QA qubits is sparse, thereby limiting users to only specify J_{ij} for those qubits that are physically connected. Thus, the abstract problem Hamiltonian is *embedded* into QA hardware where a program qubit (Q_i) with higher connectivity is represented by multiple physical qubits (q_i) called *chain* (Fig. 3). Satisfying the following conditions is sufficient to guarantee that both the abstract Hamiltonian and the embedded Hamiltonian executed on the QA hardware have identical ground states:

1. All chains representing program qubits must be a connected component graph—i.e., there must be a path between any two qubits within a chain.
2. There must be at least one connection between chains whose corresponding program qubits are connected.
3. The quadratic coefficient J_{ij} is distributed equally among the couplers connecting Q_i and Q_j .
4. The linear coefficient $\mathbf{h}_i$ is distributed equally among all physical qubits of the corresponding chain.
5. Inter-chain quadratic coefficients must be large enough to guarantee that all qubits within a chain take an identical value—i.e., a very high penalty for broken chains.

2.6 Prior Work Limitations

2.6.1 Circuit Cutting in Digital QCs. Circuit cutting techniques, namely CutQC [87], partition quantum circuits into smaller sub-circuits, enabling larger quantum circuits to be run on smaller QCs. However, a similar approach is infeasible in the analog quantum realm because: (a) analog QAs do not incorporate quantum circuits to cut its wires; and (b) partitioning graphs by edge/node removal is nontrivial (e.g., highly dense graphs are non-partitionable).

2.6.2 Solving Larger Problems on Smaller QAs. Previous methods for solving larger problems on smaller QAs [63, 75] employ iterative or alternating approaches involving approximations, leading to reduced reliability as problem size increases. Additionally, convergence—even to a local optimum—is not guaranteed with these techniques. Conversely, Skipper explores the entire search space comprehensively without resorting to approximations, and since it is not iterative, it does not face convergence issues.

2.6.3 Application-Specific Policies. Recent studies have proposed methods for tackling larger instances in various domains, such as Boolean Satisfiability (SAT) [86], Max-Clique [71, 73, 74, 76], and compressive sensing with matrix uncertainty [18, 56]. However, these techniques are tailored to their specific applications and cannot be easily adapted to other domains. In contrast, Skipper is versatile and can be applied to any problem Hamiltonian. Moreover, reduction to SAT and Max-Clique often leads to a polynomial increase in program qubits, expanding the problem size.

2.6.4 FrozenQubits. Skipper is inspired by FrozenQubits [9], with both methods aiming to eliminate high-degree program qubits. While the impact of FrozenQubits on addressing larger problems in digital QCs is minimal due to the one-to-one correspondence between program and physical qubits, Skipper, on the other hand, is capable of solving larger problems on QAs and enhancing QA fidelity. Moreover, unlike FrozenQubits, whose performance declines with increasing graph density, Skipper maintains effectiveness across a spectrum of graph densities, from sparse to dense structures.

2.7 Goal of This Paper

Figure 4(a) shows the maximum and average chain lengths for different graph topologies embedded on a 5761-qubit QA. A few dominant chains contain over 7.9x as many qubits as the average chain lengths. Furthermore, Fig. 4(b) displays the number of unused qubits when embedding the largest possible graphs on a 5761-qubit QA for different graph topologies, indicating that more than 25% of physical qubits remain unutilized, primarily due to dominant chains.

The severe underutilization of QA qubits, along with utilizing several physical qubits to represent a single program qubit, severely diminishes the capacity of QAs by up to 33x. For instance, while current D-Wave QAs boast over 5,700 qubits, they can accommodate at most 177 program qubits with full connectivity. The aim of this paper is to enable QAs to tackle larger problems by pruning dominant chains, while also enhancing the fidelity of the QAs.




Figure 4. Maximum embeddable BA graphs on 5761-qubit QA: (a) Avg and Max chain lengths, and (b) Number of unutilized qubits.

3 Methodology

3.1 Hardware Platform

For our evaluations, we utilize the D-Wave Advantage System (version 6.2), which features over 5,760 qubits and more than 40,000 couplers, accessed via the D-Wave Leap cloud service [39]. We employ the default annealing time of 20 microseconds and adhere to the anneal schedule recommended for this device. Each problem is run for 4,000 trials to comply with the two seconds maximum job duration limit.

3.2 Software Platform

We utilize the *minorminer* tool [23, 85] to find embeddings for arbitrary problem Hamiltonians on QA hardware. In our experiments, we set a timeout of 1,000 seconds, a maximum of 20 failed attempts for improvement, and conduct 20 trials. To program the D-Wave QAs, we employ the Ocean SDK [84].

3.3 Benchmarking

Although current QAs feature over 5,700 qubits, their single-instruction operation model limits them to a few hundred program qubits with higher degrees, which is far below the number of variables required for real-world applications. Consequently, in this study, we employ synthetic benchmarks instead of real-world problems. In many real-world applications, graphs often exhibit a “Power-Law” distribution [1, 25, 30, 31, 37, 54, 66], and the *Barabasi–Albert* (BA) algorithm [5, 19] is considered representative of these real-world graph structures [9, 20, 34, 43, 47, 97, 101, 102]. The BA graphs are generated with a preferential attachment factor m , enabling us to vary the density of the graphs by adjusting m —with higher values of m yielding denser graphs. We generate BA graphs with m values ranging from $m = 1$ (BA-1) to $m = 6$ (BA-6) to capture a broad spectrum of topologies, from sparse to nearly fully connected networks, thus effectively representing the dynamics of various real-world systems [25]. Edge weights are assigned randomly following a standard normal distribution, which is a common approach in QA benchmarking [10, 11, 27].

3.4 Figure of merit

In our evaluations, we use the *Energy Residual* (ER) to assess the fidelity of QA as

$$\downarrow \text{Energy Residual (ER)} = |E_{\min} - E_{\text{global}}|, \quad (2)$$

where E_{global} represents the global minimum of the benchmark problem, and $E_{\min}$ corresponds to the best solution obtained by the QA. Ideally, an ER value closer to zero is desirable as it indicates a solution that closely aligns with the ground state of the problem Hamiltonian. We conducted intensive classical computations using state-of-the-art tools [12] to determine the global minima of the benchmarks.

4 Skipper: Skipping Dominant Chains

We introduce *Skipper*, a software technique to enhance the capacity and fidelity of QAs through strategically skipping dominant qubit chains.

4.1 Key Insights

4.1.1 Not All Program Qubits are Equal. In digital QCs, the individual fabrication of physical qubits, such as superconducting ones, results in inevitable performance variations [90]. Compilers, therefore, aim to prioritizing high-quality ones and limit the reliance on those of lower quality [46, 58, 90]. However, in analog QCs, our observations reveal a significant variability at the level of program qubits.

Figure 5(a) shows the histogram of chain lengths (in log-scale) for the BA-3 graph type after embedding onto a 5761-qubit QA device, revealing a *Power-Law* distribution with some notably longer *dominant chains* and a majority of considerably shorter chains. Figure 5(b) presents the maximum and average chain lengths as the number of nodes in BA-3 graphs increases, notably magnifying the variability in chain lengths with the increase in problem size.

These intriguing observations extends beyond the BA-3 graph type, and we observe it in all benchmark graphs, including BA-1 to BA-6, spanning from sparse to nearly fully connected graphs. Additionally, we observe the nonuniformity of chain lengths in regular and fully connected graphs.




Figure 5. (a) Histogram of chain lengths for a 600-node BA-3 graph (log-scale), indicating a “Power-Law” distribution of chain lengths. (b) Max and Avg chain lengths of BA-3 graphs, embedded on a 5761-qubit QA.

4.1.2 QA Qubits are significantly Underutilized. We observe that, on average, 25% of physical qubits remain unused as they get trapped by chains. Additionally, we observe that the dominant chains significantly contribute to this qubit isolation. This underutilization of QA qubits, along with utilizing several physical qubits to represent a single program qubit, severely diminishes the capacity of QAs by up to 33x. For instance, the 2048-qubit and 5760-qubit QAs by D-Wave can accommodate a maximum of 64 and 177 fully connected program qubits, respectively.

4.1.3 Diminishing Returns with Increased QA Trials.

QAs are noisy and prone to errors, leading to a systematic bias during the execution of quantum programs. This bias causes deviations from the global optimum, reducing the reliability of QAs [4, 10, 11, 51]. The bias arises from repeating the same quantum program across multiple iterations, exposing all trials to a similar noise profile [10].

Figure 6 shows that when the number of trials in QA is increased, the output distribution reaches saturation. This indicates that the gap between the ideal solution and the QA does not reduce despite drawing more samples. Moreover, due to the operation model of QAs as single-instruction computers, strategies commonly used in gate-based QCs [67, 89] to address this bias are not applicable.




Figure 6. The Energy Residual (ER) in QAs tends to plateau with an increasing number of trials, and the global minimum often remains unreachable by QAs.

4.2 Overview of Skipper

Figure 7 shows the overview of Skipper. Skipper leverages insights into the distribution of chain lengths and the severe underutilization of qubits in QAs, employing a strategic approach to prune dominant chains and replace the corresponding program qubit with two potential measurement outcomes (+1 and -1). This process involves eliminating each chain, which partitions the search space into two sub-spaces. Skipper explores all sub-spaces, guaranteeing an exact recovery of the optimum solutions.

Eliminating a dominant chain accomplishes two significant objectives: firstly, it frees up physical qubits previously used within pruned chains, and secondly, it eliminates the isolation of solitary qubits resulting from dominant chains. As a result, Skipper enables the handling of larger problems by accommodating a significantly higher number of program qubits. Additionally, Skipper significantly enhances QA fidelity by substantially mitigating the impact of dominant chains, a primary factor in compromising QA reliability. While Skipper utilizes more quantum resources due to the need to execute 2^c unique quantum programs for the removal of c chains, it doesn’t correspondingly enhance QAs (baseline) performance, as demonstrated in Fig. 6.




Figure 7. Overview of Skipper.

4.3 Chain Skips: How, Where, and When to Skip?

Figure 8 illustrates the elimination of two chains from a problem with five variables, creating two and four *independent* sub-problems, respectively. To skip a chain, the program qubit is replaced with +1 and -1 (two measurement outcomes), removing the node and its connected edges from the problem graph. Unlike digital QCs, where removing one program qubit results in reducing the physical qubit utilization by one, in QAs, removing one program qubit liberates all the physical qubits involved in its corresponding chain.




Figure 8. By replacing Q_0 with +1 and -1 among five spin variables (baseline), two sub-problems each with four spin variables are obtained ($c = 1$). Fixing Q_1 in these two sub-problems with +1 and -1 results in four sub-problems with three spin variables ($c = 2$). The same embedding is utilized for all sub-problems at each level of the binary tree.

Identifying dominant chains to trim in Skipper is nontrivial. In digital QCs, high-degree program qubits necessitate more CNOT gates, enabling direct identification prior to circuit compilation. However, in QAs, it is not feasible to directly recognize program qubits linked to longer chains, thus requiring embedding techniques to identify them. Furthermore, it is not always optimal to prune the dominant chain. In Fig. 9(a), the dominant chain is Q_0 and consists of ten physical qubits. Pruning Q_0 (Fig. 9(b)), liberates all ten physical qubits, leaving the other chains intact. However, as shown in Fig. 9(c), removing Q_2 and re-embedding the problem not only releases the five physical qubits associated with Q_2 but also effectively reduces Q_0 to a singleton chain, totaling fourteen qubits released.

Skipper adopts a greedy approach to prune c chains by sorting program qubits based on their degree and removing the top c qubits simultaneously. The removal of a single program qubit can have a substantial impact on other chains, as shown in Fig. 9(c). In the context of irregular graphs that often follow the Power-Law distribution in real-world applications, this greedy approach exhibits a desirable, near-optimal behavior for $c \geq 5$ chain cuts.




Figure 9. (a) Embedding of five program qubits on a grid. (b) Freeing ten qubits by pruning the dominant chain Q_0 . (c) Fourteen qubits freed by pruning Q_2 .

4.4 Skip Count: A Cost-Performance Tradeoff

Skipper permits users to trim up to eleven chains. Each chain skipped bifurcates the search space; therefore, trimming up to eleven chains can lead to a maximum of 2048 sub-problems. Skipper executes all corresponding QMIs to ensure exact solution recovery. However, the nontrivial embedding process and the need to execute up to 2048 embeddings can create a bottleneck for Skipper. Fortunately, the identical structure of all sub-problems at the c -th level in the binary tree enables sharing the same embedding across them (Fig. 8).

4.5 Unembedding: Remediating Broken Chains

The sub-problem resulting from cutting chains consists of $n - c$ program qubits. However, after embedding, the problem executed on the QA hardware encompasses N physical qubits, where $c \ll n \ll N$. As a result, the QA produces outcomes as N -bit strings, with each program qubit collectively represented by multiple bits in a chain. Therefore, Skipper *unembeds* these outcome samples, converting them back into the space of program qubits.

Ideally, all physical qubits within a chain should have identical values in a given QA sample. The value of the associated program qubit is then determined by observing any one of the physical qubits within it (e.g., program qubit Q_0 in Fig. 10). However, QAs are inherently open systems, as interactions with the environment are unavoidable in QCs, and the annealing process tends to be diabatic since truly adiabatic processes are often unfeasible [11]. As a result, qubits within a chain can take different values, an issue known as *broken chains* [21, 33, 44, 72].

To remediate broken chains, Skipper employs the *majority voting* approach. For instance, in Fig. 10, although Q_1 exhibits a broken chain with varying qubit values, the unembedding process assigns a value of -1, reflecting the majority of -1 values within the chain (4 versus 1). However, not all chains have an odd length, and forcing the embedding to produce odd chain lengths is nontrivial. Unembedding even length chains with mostly identical qubit values (e.g., Q_2 in Fig. 10) is not challenging, as majority voting can effectively determine the value of the program qubit. However, as demonstrated by Q_3 in Fig. 10, a chain of even length can contain an equal number of -1 and +1 values, referred to as *balanced chains*, a condition where majority voting fails. Skipper manages balanced chains by counting them and implementing distinct strategies based on their quantity. For problems with fewer than ten balanced chains, Skipper discards their qubit values and uses a brute-force approach (with up to 1024 possible configurations), selecting the configuration that yields the lowest energy value. If the number of balanced chains exceeds ten, Skipper randomly assigns values to the corresponding program qubits. When a broken chain occurs, Skipper can optionally apply Single-Qubit Correction (SQC) [10, 11] postprocessing to maintain a feasible solution for the original problem.




Figure 10. Unembedding examples

4.6 Decoding Sub-Problem Results

After unembedding, each sample will encompass $n - c$ bits, while the original problem includes n variables. The decoding process reintroduces the values of the c pruned program qubits, which were fixed during the sub-problem formulation by assigning fixed values to these variables.

4.7 Postprocessing

Theoretically, QAs sample from a Boltzmann distribution, exponentially favoring lower energy values, and thus should locate the global optimum in few attempts. However, like other QCs, QAs are vulnerable to noise and various error sources that degrade their fidelity. To enhance the reliability of QAs, we can optionally apply postprocessing heuristics to the resulting samples [11].

4.8 Deriving the Final Output

In Skipper, all sub-problems are executed independently, each one corresponding to a separate sub-space of the primary problem. Consequently, in Skipper, the sample with the lowest energy or objective value is deemed as the ultimate output, with the originating sub-space of this global optimum being of no consequence.

4.9 Overhead of Skipper

Let c represent the number of skipped chains, e denote the edges in the problem graph, r symbolize the number of trials on the QA, while n and N correspond to the number of program and physical qubits, respectively.

Quantum overhead: Skipper allows for up to eleven chain cuts, necessitating the execution of at most 2048 distinct quantum executables, each running independently.

Classical overhead: We separate the embedding overhead of Skipper from all other classical pre/post-processing modules, as the embedding is the primary factor influencing the end-to-end runtime of the proposed techniques in this paper (refer to section 8). Given the fact that $c \ll n \ll N \ll r$, Skipper demonstrates a classical time complexity of $O(2^c(rN + c))$. This is representative of the unembedding and decoding processes for outcome samples of sub-problems.

Embedding overhead: In Skipper, all sub-problems of the binary tree at the c -th level share a single embedding, leading to $O(1)$ embedding complexity. Note that we assume all sub-problems are executed on the same QA hardware or that all devices have the same working graph topology, allowing them to share the embedding.

Memory utilization: The memory utilization in Skipper scales according to $O(rN2^c)$.

5 Skipper Evaluation Results

We evaluate Skipper using Barabasi–Albert (BA) graphs [19] with different preferential attachment factor values: $m = 1$ (BA-1) to $m = 6$ (BA-6).

5.1 Solving Larger Problems

5.1.1 Impact on Chain Length. Figure 11(a) illustrates that increasing the number of chain cuts (c) in Skipper leads to a reduction in the average chain length of the embeddings. Figure 11(b) demonstrates that Skipper decreases the mean chain length by up to 1.32x (with an average of 1.22x) when cutting up to eleven chains.

Figure 12(a) shows that the maximum chain length of the embeddings decreases as c in Skipper increases. Figure 12(b) shows that cutting up to $c = 11$ chains in Skipper reduces the maximum chain length by up to 9.14x (average 1.86x). Our observations indicate that long chains are the primary contributing factor to the underutilization of physical qubits.




Figure 11. Relative Avg. chain length in Skipper compared to baseline (lower is better). (a) Relative Avg. chain length for different graphs as cut size (c) increase. (b) Overall relative mean chain lengths for up to 11 chain cuts.

5.1.2 Impact on Qubit Utilization. Figure 13(a) displays the average and maximum number of physical qubits when up to eleven chains are pruned. In Fig. 13(b), Skipper reduces underutilization of QA qubits by up to 57% (average 22.14%) with up to eleven trimmed chains.

5.1.3 Impact on Capacity of QAs. The QA capacity to accommodate specific graph types, from BA-1 to BA-6, is determined by the largest number of program qubits of each type that can be embedded on the QA. Figure 14(a) shows that QA capacity in Skipper improves with increasing c across various graph topologies.

Figure 14(b) demonstrates that Skipper enables the embedding of larger problems onto current QAs, with an increase of up to 59.61% (average 28.26%). It is important to note that this growth in the number of program qubits necessitates a substantial increase in the number of physical qubits, as one program qubit is represented by multiple physical qubits.




Figure 12. Relative Max chain length in Skipper compared to the baseline (lower is better). (a) Relative Max chain length for different graphs as c increases. (b) Overall relative max chain lengths for up to 11 chain cuts.




Figure 13. (a) Utilization of Physical Qubits in Skipper across Different Graph Types. (b) Relative Number of Unused Physical Qubits in Skipper for up to 11 Chain Cuts, Compared to the Baseline. Lower is better.




Figure 14. Relative QA capacity in Skipper compared to baseline. (a) Relative capacity for different graphs as cuts increase. (b) Overall relative capacity for up to 11 chain cuts. Higher is better.

Skipper's performance remains consistent regardless of the increasing density of problem graphs (from BA2 to BA6).

5.2 Boosting QA Reliability

In addition to enhancing QA capacity, Skipper can be employed to improve the reliability of QAs.

5.2.1 Impact on Embedding Quality. QAs do not incorporate circuits, thus precluding the use of the Probability of Successful Trials metric commonly employed to assess compilation quality in digital QCs [3, 9, 62, 88]. Prior studies suggest that embeddings with similar chain lengths can produce better solutions [22, 24, 81, 92]. Figure 15(a) demonstrates that trimming up to eleven chains in Skipper reduces the average variance in chain lengths by 2.93x (up to 70.19x).




Figure 15. (a) Relative variance of chain lengths and (b) relative embedding time in Skipper compared to the baseline when trimming up to eleven chains. Lower is better.

5.2.2 Impact on Embedding Time. Figure 15(b) demonstrates that pruning up to eleven chains in Skipper leads to a significant reduction in embedding time, with a maximum improvement of 17.13x (average improvement of 7.12x).

5.2.3 Impact on Fidelity. Figure 16(a) shows that as Skipper skips more chains, the Energy Residual (ER) decreases, indicating a progressive approach towards the global optimum. Additionally, Fig. 16(b) demonstrates a significant reduction in ER by up to 44.4% (average 33.08%), when up to five chains are cut using Skipper, compared to the baseline.




Figure 16. Relative Energy Residual (ER) in Skipper compared to baseline (lower is better). (a) Relative ER as c increase. (b) Overall relative ER for up to five chain cuts.

6 Skipper-G: A Markovian Approach

We propose *Skipper-G*, a greedy scheme that reduces the quantum cost of Skipper by skipping the examination of sub-problems unlikely to contain the global optimum. However, this strategy entails a trade-off: Skipper-G achieves marginally lower fidelity gains compared to Skipper and is ineffective for enhancing QA capacity.

6.1 Insight: Not ALL Sub-Spaces Include Global Optimum

Skipper employs a Breadth-First Search (BFS) strategy to examine sub-problems, as depicted in Fig. 17(a). Trimming each chain bifurcates the search space, with skipping c chains resulting in a binary tree of depth c . To ensure successful recovery, Skipper evaluates all leaf nodes, running a separate QMI for each sub-space at the tree's last level. Notably, Skipper does not examine intermediate nodes (or sub-spaces) since all chains are trimmed simultaneously.

Users define the number of chain cuts in Skipper, with the option to skip up to eleven chains based on their budgetary constraints. For instance, if a user opts for the maximum allowable eleven cuts, Skipper must run 1024 QMIs when all linear coefficients are zero [9], and up to 2048 QMIs otherwise. Nonetheless, these sub-problems are independent, allowing for parallel execution by Skipper. Notably, Skipper's overall runtime remains comparable to the baseline, attributed to the significantly reduced embedding time, as detailed in Section 5.2. However, the quantum costs incurred on QCs are substantially higher than those on classical platforms, which may present affordability issues for some users.

Not every sub-space contains the global optimum. Leveraging this insight, we introduce *Skipper-G* (*greedy Skipper*), which reduces the quantum cost of Skipper by adopting a Depth-First Search (DFS) strategy (Fig. 17(b)), to bypass sub-spaces unlikely to include the global optimum. When pruning the maximum of eleven chains, Skipper-G executes 23 QMIs, in contrast to Skipper's potential 2048 QMIs.




Figure 17. (a) Skipper utilizes a Breadth-First Search (BFS) strategy, examining all leaf nodes (intermediate nodes are not examined). (b) Skipper-G adopts a Depth-First Search (DFS) strategy, examining only two nodes at each level of the binary tree (including intermediate nodes). Example: With $c = 11$ chain cuts, Skipper and Skipper-G execute at most 2048 and 23 QMIs, respectively.

6.2 How Skipper-G Work?

Figure 18 illustrates the overview of the Skipper-G scheme. In Skipper-G, similar to Skipper, users can determine the number of chain cuts, with the possibility of skipping up to eleven chains, depending on their budget constraints. However, unlike Skipper where all chains are cut simultaneously, Skipper-G employs an iterative approach, cutting one chain in each iteration. As illustrated in Fig. 17(b), Skipper-G initiates by setting the root node (i.e., the baseline with no chains cut) as the current node and executing the corresponding quantum program. For each chain cut, Skipper-G performs the following steps:

1. In the current node (problem), the dominant chain is trimmed by setting its corresponding program qubit to either +1 or -1, resulting in two child nodes. If the current node at level c has the index x , then its left and right children at level $c + 1$ will have indices $2x$ and $2x + 1$, respectively (e.g., node $x = 1$ at the third level leads to nodes 2 and 3 in Fig. 17(b)).
2. The quantum programs corresponding to the children are executed on the QA.
3. The best offspring is set as the current node.

6.3 Branch and Bound Criteria

When evaluating a node in Skipper-G, a quantum program is executed on a QA device for multiple trials. Each trial produces an outcome with an associated objective value. The assessment of node quality in Skipper-G is based on the following feature (lower is better):

$$\downarrow f(Z) = \left| \frac{1}{E_{\min} \times EV} \right|, \quad (3)$$

where Z denotes the set of obtained samples, and $E_{\min}$ and EV represent the minimum and the expected value of the energy values in Z , respectively. The lower the value of f , the greater the likelihood that a child includes the global optimum in its corresponding subspace during the traversal of the associated binary tree. This feature balances the best sample with the overall quality of all samples, reducing the likelihood of getting trapped in local optima.

6.4 Overhead of Skipper-G

Skipper-G is capable of trimming up to eleven chains, which necessitates a maximum of 23 distinct QMI executions. Although Skipper-G examines two nodes at each level of the binary tree, these nodes, due to their identical structures, can utilize a single embedding. Consequently, Skipper-G requires c embeddings for c chain cuts. Nonetheless, since these embeddings can be executed in parallel, the embedding time for Skipper-G remains similar to that of the baseline, as the root node's embedding is expected to be more time-consuming than the embeddings of the smaller intermediate nodes.




Figure 18. Overview of Skipper-G.

6.5 Evaluation Results

Figure 19(a) illustrates the ER for various chain cut counts in Skipper-G, showing that progressively trimming more dominant chains leads to a decrease in the ER, approaching the global minimum. Additionally, Fig 19(b) reveals that pruning up to five chains in Skipper-G can reduce the gap between the global optimum and the best QA sample by as much as 40.75% (Avg. 29.19%) compared to the baseline.

Skipper marginally outperforms Skipper-G, achieving a 3.89% greater reduction in ER, albeit at the expense of significantly higher quantum resource utilization. Skipper-G includes the baseline at the root of the binary tree, ensuring it performs no worse than the baseline.




Figure 19. Relative ER for five chain cuts compared to the baseline (lower is better). (a) Skipper-G: Relative ER with increasing C . (b) Overall Relative ER (lower is better) for Skipper vs. Skipper-G.

7 Skipper and Skipper-G in Classical Realm

Unfortunately, neither Skipper nor Skipper-G can be utilized to enhance the fidelity of optimization techniques used in classical realm. In the classical domain, the hardness of optimization problems depends on the number of variables and the graph topology. For instance, while planar graphs are tractable [28, 35] in classical realm, neither regular nor power-law graphs become planar simply by eliminating a few nodes. Additionally, eliminating ten nodes from a 1000-node graph results in sub-graphs with 990 nodes, which typically remain intractable in the classical realm.

Similarly, Skipper is not suitable for tackling larger problems in the classical realm. Its primary goal is to address the sparse connectivity of qubits, a key factor limiting the capacity of QAs. However, the full connectivity of classical bits does not present a similar limitation.

8 Workflow Analysis

The runtime of quantum programs is mainly determined by queuing delays, execution modes through cloud services (which vary across providers), and embedding time, rather than the execution time on quantum hardware (microseconds to milliseconds). To offer a holistic examination of the runtime between the proposed techniques and the baseline, we employ the following analytical model:

$$T = T_{\text{emb}} + N_{\text{QMI}} (T_{\text{queue}} + T_{\text{QMI}} + T_{\text{net}}) + T_{\text{classical}}, \quad (4)$$

where T_{emb} is the embedding time, N_{QMI} is the number of quantum executables, T_{queue} is the job wait time, T_{QMI} is the QMI execution time, T_{net} is the network delay, and $T_{\text{classical}}$ is the classical pre/post-processing time.

For r trials, $T_{\text{QMI}} = t_p + \Delta + r \times t_s$, where t_p is the raw signal preparation time, Δ is the 10ms QA initialization time, and t_s is the single annealing/readout time. Given that D-Wave limits T_{QMI} to two seconds, we assume $T_{\text{QMI}} = 2$ in all cases. We assume one second for T_{net} for each job.

We assume a baseline embedding time of 30 minutes, decreasing proportionally with skipped chains (as discussed in section 5.2). For example, pruning ten chains reduces the embedding time to three minutes. All embeddings can be computed in parallel, making T_{emb} in Skipper-G the maximum time for individual embeddings. Additionally, we allocate one second each for pre- and post-processing.

We examine two access scenarios: *shared* and *dedicated*, with one and zero-second queuing times, respectively. Figure 20 compares the end-to-end runtime of the baseline and our proposed techniques with $c = 11$, resulting in up to 1024 QA runs. Skipper shows significantly greater advantages over others in the dedicated access mode.




Figure 20. Overall Runtime comparison.

9 Related Work

Prior studies can be broadly classified into two categories: (a) techniques for solving larger problems on QAs, which are relevant primarily to Skipper; and (b) approaches for enhancing QA fidelity, which are considered related work to both Skipper and Skipper-G.

Prior research on solving larger problems with smaller QAs [63, 75] are iterative schemes, which tend to lose reliability as problem size increases due to reliance on approximations. Conversely, Skipper explores the entire search space

without resorting to approximations. Recent studies have introduced schemes for addressing larger instances of Boolean Satisfiability (SAT) [86], Max-Clique [71, 73, 74, 76], and compressive sensing with matrix uncertainty [18] problems. However, these methods are specific to their respective applications and are not transferable to other domains, whereas Skipper is versatile and applicable to any application.

Policies for improving the fidelity of QAs can be classified as: (a) Preprocessing [10, 17, 70], modifying QMIs before submission; (b) Postprocessing [11], enhancing outcomes using heuristics; (c) Hybrid strategies [8, 13], combining heuristics and QAs for reliability; (d) Logical analog qubits [42, 48–50, 53, 55, 68, 79, 82, 92, 94–96, 100], spreading qubit information over multiple physical qubits; and ensembling policies [10, 16, 17], subjecting the quantum program to different noise profiles to suppress the bias. These proposals are orthogonal to Skipper and Skipper-G and can effectively boost the reliability of our proposed techniques.

Skipper is inspired by FrozenQubits [9] in digital QCs. While FrozenQubits enhances fidelity of optimization applications in digital QCs, Skipper excels in addressing larger problems and enhancing QA fidelity. More importantly, while FrozenQubits’ performance diminishes with increased problem graph density, Skipper and Skipper-G maintain their performance, demonstrating the effectiveness of our proposal in handling sparse to dense graphs.

10 Conclusion

We propose *Skipper*, a software scheme designed to enhance the capacity and fidelity of QAs. Observing that chain lengths in QAs follow a “Power-Law” distribution, with a few *dominant chains* containing significantly more qubits than others, Skipper prunes these chains. This approach replaces their corresponding program qubits with two possible measurement outcomes, freeing all qubits in the dominant chains and an additional 25% of isolated qubits previously entrapped in chains. Our experiments on a 5761-qubit QA by D-Wave show that Skipper allows QAs to solve problems up to 59% larger (Avg. 28.3%) when up to eleven chains are skipped. Additionally, by removing five chains, Skipper substantially improves QA fidelity by up to 44.4% (Avg. 33.1%).

The number of chain cuts in Skipper is user-defined; users can trim up to eleven chains, which necessitates running an average of 1024 (and up to 2048) distinct quantum executables. However, this may lead to affordability concerns for some users. To mitigate this, we introduce *Skipper-G*, a greedy scheme that prioritizes examining sub-spaces more likely to contain the global optimum. When up to eleven chains are pruned, Skipper-G runs a maximum of 23 quantum executables. Our experiments show that Skipper-G enhances QA fidelity by up to 40.8% (Avg. 29.2%), requiring only 11 quantum executable runs for up to five chain cuts, compared to Skipper’s 32 runs.

References

- [1] Matthew T Agler, Jonas Ruhe, Samuel Kroll, Constanze Morhenn, Sang-Tae Kim, Detlef Weigel, and Eric M Kemen. Microbial hub taxa link host and abiotic factors to plant microbiome variation. *PLoS biology*, 14(1):e1002352, 2016.
- [2] Dorit Aharonov, Wim Van Dam, Julia Kempe, Zeph Landau, Seth Lloyd, and Oded Regev. Adiabatic quantum computation is equivalent to standard quantum computation. *SIAM review*, 50(4):755–787, 2008.
- [3] Mahabubul Alam, Abdullah Ash-Saki, and Swaroop Ghosh. Circuit compilation methodologies for quantum approximate optimization algorithm. In *MICRO 2020*. IEEE, 2020.
- [4] Tameem Albash and Daniel A Lidar. Adiabatic quantum computation. *Reviews of Modern Physics*, 90(1), 2018.
- [5] Reka Albert. Scale-free networks in cell biology. *Journal of cell science*, 118(21):4947–4957, 2005.
- [6] Amazon. Amazon Braket - Explore and experiment with quantum computing. <https://aws.amazon.com/braket/>, 2022. [Online; accessed 22-July-2021].
- [7] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779), 2019.
- [8] Ramin Ayanzadeh. *Leveraging Artificial Intelligence to Advance Problem-Solving with Quantum Annealers*. PhD thesis, University of Maryland, Baltimore County, 2020.
- [9] Ramin Ayanzadeh, Narges Alavisamani, Poulami Das, and Moinuddin Qureshi. Frozenqubits: Boosting fidelity of qaoa by skipping hotspot nodes. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, pages 311–324, 2023.
- [10] Ramin Ayanzadeh, Poulami Das, Swamit Tannu, and Moinuddin Qureshi. Equal: Improving the fidelity of quantum annealers by injecting controlled perturbations. In *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*, pages 516–527. IEEE, 2022.
- [11] Ramin Ayanzadeh, John Dorband, Milton Halem, and Tim Finin. Multi-qubit correction for quantum annealers. *Scientific Reports*, 11(1):16119, 2021.
- [12] Ramin Ayanzadeh, John Dorband, Milton Halem, and Tim Finin. Multi qubit correction (mqc) for quantum annealers, 2021. Python implementation of MQC.
- [13] Ramin Ayanzadeh, John Dorband, Milton Halem, and Tim Finin. Quantum-assisted greedy algorithms. In *IGARSS 2022-2022 IEEE International Geoscience and Remote Sensing Symposium*, pages 4911–4914. IEEE, 2022.
- [14] Ramin Ayanzadeh, Milton Halem, and Tim Finin. Solving hard sat instances with adiabatic quantum computers. In *AGU Fall Meeting Abstracts*, volume 2018, pages IN41B–27, 2018.
- [15] Ramin Ayanzadeh, Milton Halem, and Tim Finin. Sat-based compressive sensing. *arXiv preprint arXiv:1903.03650*, 2019.
- [16] Ramin Ayanzadeh, Milton Halem, and Tim Finin. An ensemble approach for compressive sensing with quantum annealers. In *IGARSS 2020*. IEEE, 2020.
- [17] Ramin Ayanzadeh, Milton Halem, and Tim Finin. Reinforcement quantum annealing: A hybrid quantum learning automata. *Scientific Reports*, 10(1), 2020.
- [18] Ramin Ayanzadeh, Seyedahmad Mousavi, Milton Halem, and Tim Finin. Quantum annealing based binary compressive sensing with matrix uncertainty. *arXiv:1901.00088*, 2019.
- [19] Albert-László Barabási and Réka Albert. Emergence of scaling in random networks. *science*, 286(5439):509–512, 1999.
- [20] Albert-László Barabási, Réka Albert, and Hawoong Jeong. Scale-free characteristics of random networks: the topology of the world-wide web. *Physica A: statistical mechanics and its applications*, 281(1-4):69–77, 2000.
- [21] Aaron Barbosa, Elijah Pelofske, Georg Hahn, and Hristo N Djidjev. Optimizing embedding-related quantum annealing parameters for reducing hardware bias. In *Parallel Architectures, Algorithms and Programming: 11th International Symposium, PAAP 2020, Shenzhen, China, December 28–30, 2020, Proceedings 11*, pages 162–173. Springer, 2021.
- [22] Tomas Boothby, Andrew D King, and Aidan Roy. Fast clique minor generation in chimera qubit connectivity graphs. *Quantum Information Processing*, 15(1), 2016.
- [23] Jun Cai, William G Macready, and Aidan Roy. A practical heuristic for finding graph minors. *arXiv:1406.2741*, 2014.
- [24] Vicky Choi. Minor-embedding in adiabatic quantum computation: I. the parameter setting problem. *Quantum Information Processing*, 7:193–209, 2008.
- [25] Aaron Clauset, Ellen Tucker, and Matthias Sainz. The colorado index of complex networks. *Retrieved July*, 20(2018):22, 2016.
- [26] International Business Machines Corporation. Universal Quantum Computer Development at IBM. <http://research.ibm.com/ibm-q/research/>, 2021. [Online; accessed 22-July-2021].
- [27] Arnab Das and Bikas K Chakrabarti. Colloquium: Quantum annealing and analog quantum computation. *Reviews of Modern Physics*, 80(3), 2008.
- [28] Vladimir G Dei, Bettina Klinz, Gerhard J Woeginger, et al. Exact algorithms for the hamiltonian cycle problem in planar graphs. *Operations Research Letters*, 34(3):269–274, 2006.
- [29] Nada Elsokkary, Faisal Shah Khan, Davide La Torre, Travis S Humble, and Joel Gottlieb. Financial portfolio management using d-wave quantum optimizer: The case of abu dhabi securities exchange. Technical report, Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States), 2017.
- [30] D Gamermann, J Triana-Dopico, and R Jaime. A comprehensive statistical study of metabolic and protein–protein interaction network properties. *Physica A: Statistical Mechanics and its Applications*, 534:122204, 2019.
- [31] Kwang-Il Goh, Eulsik Oh, Hawoong Jeong, Byungnam Kahng, and Doochul Kim. Classification of scale-free networks. *Proceedings of the National Academy of Sciences*, 99(20):12583–12588, 2002.
- [32] Google. Google Quantum AI. <https://quantumai.google/>, 2022. [Online; accessed 22-July-2021].
- [33] Erica Grant and Travis S Humble. Benchmarking embedded chain breaking in quantum annealing. *Quantum Science and Technology*, 7(2):025029, 2022.
- [34] Caitlin Gray, Lewis Mitchell, and Matthew Roughan. Super-blockers and the effect of network structure on information cascades. In *Companion Proceedings of the The Web Conference 2018*, pages 1435–1441, 2018.
- [35] Frank Hadlock. Finding a maximum cut of a planar graph in polynomial time. *SIAM Journal on Computing*, 4(3):221–225, 1975.
- [36] Honeywell. Honeywell Quantum Solutions. <https://www.honeywell.com/us/en/company/quantum>, 2021. [Online; accessed 22-July-2021].
- [37] Thomas House, Jonathan M Read, Leon Danon, and Matthew J Keeling. Testing the hypothesis of preferential attachment in social network formation. *EPJ Data Science*, 4(1):1–13, 2015.
- [38] Feng Hu, Lucas Lamata, Mikel Sanz, Xi Chen, Xingyuan Chen, Chao Wang, and Enrique Solano. Quantum computing cryptography: Finding cryptographic boolean functions with quantum annealing by a 2000 qubit d-wave quantum computer. *Physics Letters A*, 384(10), 2020.
- [39] D-Wave Systems Inc. The first and only quantum computer built for business. <https://www.dwavesys.com/>, 2022. [Online; accessed 22-July-2021].

- [40] Daisuke Inoue, Akihisa Okada, Tadayoshi Matsumori, Kazuyuki Aihara, and Hiroaki Yoshida. Traffic signal optimization on a square lattice with quantum annealing. *Scientific reports*, 11(1), 2021.
- [41] IonQ. Ionq - trapped ion quantum computing. <https://ionq.com/>, 2023. [Online; accessed 01-June-2023].
- [42] Stephen P Jordan, Edward Farhi, and Peter W Shor. Error-correcting codes for adiabatic quantum computation. *Physical Review A*, 74(5), 2006.
- [43] Sung-Soo Kim, Young-Min Kang, and Young-Kuk Kimt. Sparsity-aware reachability computation for massive graphs. In *2022 IEEE International Conference on Big Data and Smart Computing (BigComp)*, pages 157–160. IEEE, 2022.
- [44] Andrew D King and Catherine C McGeoch. Algorithm engineering for a quantum annealing platform. *arXiv preprint arXiv:1410.2628*, 2014.
- [45] Andrew D King, Jack Raymond, Trevor Lanting, Sergei V Isakov, Masoud Mohseni, Gabriel Poulin-Lamarre, Sara Ejtemaee, William Bernoudy, Isil Ozfidan, Anatoly Yu Smirnov, et al. Scaling advantage over path-integral monte carlo in quantum simulation of geometrically frustrated magnets. *Nature communications*, 12(1), 2021.
- [46] Gushu Li, Yufei Ding, and Yuan Xie. Tackling the Qubit Mapping Problem for NISQ-Era Quantum Devices. *arXiv:1809.02573*, 2018.
- [47] David Lusseau. The emergent properties of a dolphin social network. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 270(suppl_2):S186–S188, 2003.
- [48] Shunji Matsuura, Hidetoshi Nishimori, Tameem Albash, and Daniel A Lidar. Mean field analysis of quantum annealing correction. *Physical review letters*, 116(22), 2016.
- [49] Shunji Matsuura, Hidetoshi Nishimori, Walter Vinci, Tameem Albash, and Daniel A Lidar. Quantum-annealing correction at finite temperature: Ferromagnetic p-spin models. *Physical review A*, 95(2), 2017.
- [50] Shunji Matsuura, Hidetoshi Nishimori, Walter Vinci, and Daniel A Lidar. Nested quantum annealing correction at finite temperature: p-spin models. *Physical Review A*, 99(6), 2019.
- [51] Catherine C McGeoch. Theory versus practice in annealing-based quantum computing. *Theoretical Computer Science*, 2020.
- [52] Microsoft. Azure Quantum - Quantum Service | Microsoft Azure. <https://azure.microsoft.com/en-us/services/quantum/#product-overview>, 2022. [Online; accessed 22-July-2021].
- [53] Anurag Mishra, Tameem Albash, and Daniel A Lidar. Performance of two different quantum annealing correction codes. *Quantum Information Processing*, 15(2), 2016.
- [54] Alan Mislove, Massimiliano Marcon, Krishna P Gummadi, Peter Druschel, and Bobby Bhattacharjee. Measurement and analysis of online social networks. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pages 29–42, 2007.
- [55] Naeimeh Mohseni, Marek Narozniak, Alexey N Pyrkov, Valentin Ivannikov, Jonathan P Dowling, and Tim Byrnes. Error suppression in adiabatic quantum computing with qubit ensembles. *npj Quantum Information*, 7(1), 2021.
- [56] Ahmad Mousavi, Mehdi Rezaee, and Ramin Ayanzadeh. A survey on compressive sensing: Classical results and recent advancements. *arXiv preprint arXiv:1908.01014*, 2019.
- [57] Vikram Khipple Mulligan, Hans Melo, Haley Irene Merritt, Stewart Slocum, Brian D Weitzner, Andrew M Watkins, P Douglas Renfrew, Craig Pelissier, Paramjit S Arora, and Richard Bonneau. Designing peptides on a quantum computer. *bioRxiv*, 2020.
- [58] Prakash Murali, Jonathan M Baker, Ali Javadi Abhari, Frederic T Chong, and Margaret Martonosi. Noise-adaptive compiler mappings for noisy intermediate-scale quantum computers. *arXiv preprint arXiv:1901.11054*, 2019.
- [59] Prakash Murali, Jonathan M Baker, Ali Javadi-Abhari, Frederic T Chong, and Margaret Martonosi. Noise-adaptive compiler mappings for noisy intermediate-scale quantum computers. In *ASPLOS 2019*, 2019.
- [60] Michael A Nielsen and Isaac L Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2010.
- [61] Hidetoshi Nishimori and Kabuki Takada. Exponential enhancement of the efficiency of quantum annealing by non-stoquastic hamiltonians. *Frontiers in ICT*, 4, 2017.
- [62] Shin Nishio, Yulu Pan, Takahiko Satoh, Hideharu Amano, and Rodney Van Meter. Extracting success from ibm’s 20-qubit machines using error-aware compilation. *arXiv preprint arXiv:1903.10963*, 2019.
- [63] Shuntaro Okada, Masayuki Ohzeki, Masayoshi Terabe, and Shinichiro Taguchi. Improving solutions by embedding larger subproblems in a d-wave quantum annealer. *Scientific reports*, 9(1), 2019.
- [64] Daniel O’Malley, Velimir V Vesselinov, Boian S Alexandrov, and Ludmil B Alexandrov. Nonnegative/binary matrix factorization with a d-wave quantum annealer. *PloS one*, 13(12), 2018.
- [65] PASQAL. Programmable atomic arrays - pasqal. <https://www.pasqal.com/>, 2023. [Online; accessed 01-June-2023].
- [66] Romualdo Pastor-Satorras, Claudio Castellano, Piet Van Mieghem, and Alessandro Vespignani. Epidemic processes in complex networks. *Reviews of modern physics*, 87(3):925, 2015.
- [67] Tirthak Patel and Devesh Tiwari. Veritas: accurately estimating the correct output on noisy intermediate-scale quantum computers. In *SC20. IEEE*, 2020.
- [68] Adam Pearson, Anurag Mishra, Itay Hen, and Daniel A Lidar. Analog errors in quantum annealing: doom and hope. *NPJ Quantum Information*, 5(1), 2019.
- [69] Elijah Pelofske. 4-clique network minor embedding for quantum annealers. *arXiv preprint arXiv:2301.08807*, 2023.
- [70] Elijah Pelofske, Georg Hahn, and Hristo Djidjev. Optimizing the spin reversal transform on the d-wave 2000q. *arXiv:1906.10955*, 2019.
- [71] Elijah Pelofske, Georg Hahn, and Hristo Djidjev. Solving large maximum clique problems on a quantum annealer. In *Quantum Technology and Optimization Problems: First International Workshop, QTOP 2019, Munich, Germany, March 18, 2019, Proceedings 1*, pages 123–135. Springer, 2019.
- [72] Elijah Pelofske, Georg Hahn, and Hristo Djidjev. Advanced unembedding techniques for quantum annealers. In *2020 International Conference on Rebooting Computing (ICRC)*, pages 34–41. IEEE, 2020.
- [73] Elijah Pelofske, Georg Hahn, and Hristo Djidjev. Decomposition algorithms for solving np-hard problems on a quantum annealer. *Journal of Signal Processing Systems*, 93:405–420, 2021.
- [74] Elijah Pelofske, Georg Hahn, and Hristo N Djidjev. Parallel quantum annealing. *Scientific Reports*, 12(1):4499, 2022.
- [75] Elijah Pelofske, Georg Hahn, and Hristo N Djidjev. Solving larger optimization problems using parallel quantum annealing. *arXiv preprint arXiv:2205.12165*, 2022.
- [76] Elijah Pelofske, Georg Hahn, and Hristo N Djidjev. Solving larger maximum clique problems using parallel quantum annealing. *Quantum Information Processing*, 22(5):219, 2023.
- [77] WangChun Peng, BaoNan Wang, Feng Hu, YunJiang Wang, XianJin Fang, XingYuan Chen, and Chao Wang. Factoring larger integers with fewer qubits via quantum annealing with optimized parameters. *SCIENCE CHINA Physics, Mechanics & Astronomy*, 62(6), 2019.
- [78] John Preskill. Quantum computing in the nisq era and beyond. *arXiv:1801.00862*, 2018.
- [79] Kristen L Pudenz, Tameem Albash, and Daniel A Lidar. Error-corrected quantum annealing with hundreds of qubits. *Nature communications*, 5(1), 2014.
- [80] QuEra. Quantum computing with neutral atoms. <https://www.quera.com/>, 2023. [Online; accessed 01-June-2023].
- [81] Eleanor G Rieffel, Davide Venturelli, Bryan O’Gorman, Minh B Do, Elicia M Prystay, and Vadim N Smelyanskiy. A case study in programming a quantum annealer for hard operational planning problems.

- Quantum Information Processing*, 14(1), 2015.
- [82] Mohan Sarovar and Kevin C Young. Error suppression and error correction in adiabatic quantum computation: non-equilibrium dynamics. *New Journal of Physics*, 15(12), 2013.
 - [83] Juexiao Su, Tianheng Tu, and Lei He. A quantum annealing approach for boolean satisfiability problem. In *DAC 2016*. ACM, 2016.
 - [84] D-Wave Systems. D-wave ocean sdk. <https://github.com/dwavesystems/dwave-ocean-sdk>, 2023. [Online; accessed 01-June-2023].
 - [85] D-Wave Systems. minorminer. <https://github.com/dwavesystems/minorminer>, 2023. [Online; accessed 01-June-2023].
 - [86] Siwei Tan, Mingqian Yu, Andre Python, Yongheng Shang, Tingting Li, Liqiang Lu, and Jianwei Yin. Hyqsat: A hybrid approach for 3-sat problems by integrating quantum annealer with cdcl. In *2023 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pages 731–744. IEEE, 2023.
 - [87] Wei Tang, Teague Tomesh, Martin Suchara, Jeffrey Larson, and Margaret Martonosi. Cutqc: using small quantum computers for large quantum circuit evaluations. In *ASPLOS 2021*, 2021.
 - [88] Swamit S Tannu, Poulami Das, Ramin Ayanzadeh, and Moinuddin K Qureshi. Hammer: Boosting fidelity of noisy quantum circuits by exploiting hamming behavior of erroneous outcomes. In *ASPLOS 2022*, 2022.
 - [89] Swamit S Tannu and Moinuddin Qureshi. Ensemble of diverse mappings: Improving reliability of quantum computers by orchestrating dissimilar mistakes. In *MICRO 2019*, 2019.
 - [90] Swamit S Tannu and Moinuddin K Qureshi. Not all qubits are created equal: a case for variability-aware policies for nisc-era quantum computers. In *ASPLOS 2019*, 2019.
 - [91] Tony T Tran, Minh Do, Eleanor G Rieffel, Jeremy Frank, Zhihui Wang, Bryan O’Gorman, Davide Venturelli, and J Christopher Beck. A hybrid quantum-classical approach to solving scheduling problems. In *SOCS 2016*, 2016.
 - [92] Davide Venturelli, Dominic JJ Marchand, and Galo Rojo. Quantum annealing implementation of job-shop scheduling. *arXiv:1506.08479*, 2015.
 - [93] Benjamin Villalonga, Dmitry Lyakh, Sergio Boixo, Hartmut Neven, Travis S Humble, Rupak Biswas, Eleanor G Rieffel, Alan Ho, and Salvatore Mandrà. Establishing the quantum supremacy frontier with a 281 pflop/s simulation. *Quantum Science and Technology*, 5(3), 2020.
 - [94] Walter Vinci, Tameem Albash, and Daniel A Lidar. Nested quantum annealing correction. *npj Quantum Information*, 2(1), 2016.
 - [95] Walter Vinci, Tameem Albash, Gerardo Paz-Silva, Itay Hen, and Daniel A Lidar. Quantum annealing correction with minor embedding. *Physical Review A*, 92(4), 2015.
 - [96] Walter Vinci and Daniel A Lidar. Scalable effective-temperature reduction for quantum annealers via nested quantum annealing correction. *Physical Review A*, 97(2), 2018.
 - [97] Hengbin Wang. *Complex Web-API Network Construction Based on Barabasi-Albert Model and Popularity-similarity Optimization Model*. PhD thesis, Auckland University of Technology, 2019.
 - [98] Yulin Wu, Wan-Su Bao, Sirui Cao, Fusheng Chen, Ming-Cheng Chen, Xiawei Chen, Tung-Hsun Chung, Hui Deng, Yajie Du, Daojin Fan, et al. Strong quantum computational advantage using a superconducting quantum processor. *arXiv:2106.14734*, 2021.
 - [99] Bin Yan and Nikolai A Sinitsyn. Analytical solution for nonadiabatic quantum annealing to arbitrary ising spin hamiltonian. *Nature Communications*, 13(1):2212, 2022.
 - [100] Kevin C Young, Mohan Sarovar, and Robin Blume-Kohout. Error suppression and error correction in adiabatic quantum computation: Techniques and challenges. *Physical Review X*, 3(4), 2013.
 - [101] Vladimir Nikolaevich Zadorozhnyi and Evgenii Borisovich Yudin. Structural properties of the scale-free barabasi-albert graph. *Automation and Remote Control*, 73(4):702–716, 2012.
 - [102] Stefanie Zbinden, Andreas Bärtschi, Hristo Djidjev, and Stephan Eidenbenz. Embedding algorithms for quantum annealers with chimera and pegasus connection topologies. In *High Performance Computing: 35th International Conference, ISC High Performance 2020, Frankfurt/Main, Germany, June 22–25, 2020, Proceedings*, pages 187–206. Springer, 2020.
 - [103] Alwin Zulehner, Alexandru Paler, and Robert Wille. An efficient methodology for mapping quantum circuits to the ibm qx architectures. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 38(7), 2018.