



PDF Download
3779212.3790177.pdf
30 March 2026
Total Citations: 0
Total Downloads: 445

Latest updates: <https://dl.acm.org/doi/10.1145/3779212.3790177>

RESEARCH-ARTICLE

iSwitch: QEC on Demand via In-Situ Encoding of Bare Qubits for Ion Trap Architectures

KEYI YIN, University of California, San Diego, San Diego, CA, United States

XIANG FANG, University of California, San Diego, San Diego, CA, United States

ZHUO CHEN, University of California, San Diego, San Diego, CA, United States

DAVID HAYES, Quantinuum Ltd., Cambridge, Cambridgeshire, U.K.

ENEET KAUR, Cisco Systems, San Jose, CA, United States

REZA NEJABATI, Cisco Systems, San Jose, CA, United States

[View all](#)

Open Access Support provided by:

[University of California, Los Angeles](#)

[University of California, San Diego](#)

[Cisco Systems](#)

[Quantinuum Ltd.](#)

[University of California, Berkeley](#)

[Oak Ridge National Laboratory](#)

Published: 22 March 2026

[Citation in BibTeX format](#)

ASPLOS '26: 31st ACM International
Conference on Architectural Support for
Programming Languages and Operating
Systems

March 22 - 26, 2026
PA, Pittsburgh, USA

Conference Sponsors:

SIGOPS
SIGPLAN
SIGARCH

iSwitch: QEC on Demand via In-Situ Encoding of Bare Qubits for Ion Trap Architectures

Keyi Yin

University of California, San Diego
San Diego, California, USA
keyin@ucsd.edu

David Hayes

Quantinuum
Broomfield, Colorado, USA
david.hayes@quantinuum.com

Hartmut Haeffner

University of California, Berkeley
Berkeley, California, USA
hhaeffner@berkeley.edu

Jens Palsberg

University of California, Los Angeles
Los Angeles, California, USA
palsberg@cs.ucla.edu

Xiang Fang

University of California, San Diego
San Diego, California, USA
x8fang@ucsd.edu

Eneet Kaur

Cisco Quantum Lab
San Jose, California, USA
ekaur@cisco.com

Wes Campbell

University of California, Los Angeles
Los Angeles, California, USA
wes@physics.ucla.edu

Travis Humble

Oak Ridge National Laboratory
Oak Ridge, Tennessee, USA
humblets@ornl.gov

Zhuo Chen

University of California, San Diego
San Diego, California, USA
zhc086@ucsd.edu

Reza Nejabati

Cisco Quantum Lab
San Jose, California, USA
rnejabat@cisco.com

Eric Hudson

University of California, Los Angeles
Los Angeles, California, USA
eric.hudson@ucla.edu

Yufei Ding

University of California, San Diego
San Diego, California, USA
yufeiding@ucsd.edu

Abstract

Recent advances in quantum hardware and error correction have paved the way for early fault-tolerant (EFT) quantum computing. We propose iSwitch, a hybrid system architecture for trapped-ion quantum computers (TIQC) that exploits ultra-high-fidelity single-qubit gates and efficient logical CNOTs enabled by ion shuttling. iSwitch employs bare qubits for single-qubit operations and QEC-encoded logical qubits for two-qubit gates, avoiding full logical encoding, gate synthesis, and magic state distillation. To enable this selective encoding, we develop a low-noise conversion protocol between bare and logical qubits, a hybrid instruction set tailored to two-dimensional (2D) TIQC layouts, and a compiler that minimizes conversion overhead and optimizes scheduling. Evaluations on variational quantum algorithm benchmarks show that iSwitch achieves comparable fidelity to conventional QEC methods, while reducing qubit and operation counts by roughly 33–50%, offering a practical, resource-efficient path toward EFT quantum computing.

CCS Concepts: • Hardware → Quantum error correction and fault tolerance.



This work is licensed under a Creative Commons Attribution 4.0 International License.

ASPLOS '26, Pittsburgh, PA, USA

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2359-9/2026/03

<https://doi.org/10.1145/3779212.3790177>

Keywords: Quantum computing, Quantum error correction, Ion Trap

ACM Reference Format:

Keyi Yin, Xiang Fang, Zhuo Chen, David Hayes, Eneet Kaur, Reza Nejabati, Hartmut Haeffner, Wes Campbell, Eric Hudson, Jens Palsberg, Travis Humble, and Yufei Ding. 2026. iSwitch: QEC on Demand via In-Situ Encoding of Bare Qubits for Ion Trap Architectures. In *Proceedings of the 31st ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2 (ASPLOS '26)*, March 22–26, 2026, Pittsburgh, PA, USA. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3779212.3790177>

1 Introduction

Quantum noise remains a fundamental barrier for today's quantum computing [1–3]. Various noise channels accumulate across circuit operations [4, 5], making the execution of large-scale quantum programs still practically infeasible, although modern quantum devices now already host hundreds to thousands of physical qubits across diverse hardware platforms [1, 6–10]. Programs that leave more qubits unprotected are limited to shallower circuit depths [11, 12]. This accumulation constrains the practical benefits of larger devices, showing that simply increasing qubit count does not extend computational capability.

Quantum error correction (QEC) [13–17] is widely regarded as the only viable path to overcome this bottleneck and enable fault-tolerant quantum computing (FTQC) [18–21], which is essential for solving large-scale, practically

useful problems [22–25]. With sufficient resources to implement QEC, a fault-tolerant quantum computer can in principle execute arbitrarily long circuits while maintaining high overall fidelity [4, 26].




Figure 1. Overview of iSwitch.

However, current fault-tolerant schemes, which require every operation to be executed fault-tolerantly, are extremely resource-intensive [4, 26]. The resource overhead primarily arises from two sources. First, each logical qubit must be encoded into a sufficient number of physical qubits to ensure its protection during execution. In particular, commonly used surface code encodings [16] are highly resource-intensive, requiring many physical qubits per logical qubit to achieve practical error rates. Second, the Eastin-Knill theorem [27] implies that additional resources are necessary to implement gates not native to the QEC code, such as the non-Clifford T gate in the surface code. The standard approach for realizing these non-Clifford operations, *magic state distillation* (MSD) [28–30], consumes a large number of ancilla qubits and significant runtime [31]. Overall, a fully FT design may require on the order of millions of physical qubits.

Unsurprisingly, the resource demands of full FT schemes far exceed near-term device capabilities. As noted in a recent talk by John Preskill [32], the next generation of quantum machines, dubbed Megaquop Quantum Computers (MQCs), is expected to host thousands of physical qubits, offering far greater scale than current noisy intermediate-scale quantum (NISQ) devices [33] while still falling short of full fault-tolerant capability. This intermediate regime motivates early fault-tolerant (EFT) strategies [34–36], which aim to move beyond NISQ and enable practical quantum applications by incorporating some QEC.

Among these EFT strategies, *state injection* [34, 36, 37] has emerged as a particularly promising approach. By implementing non-Clifford operations such as R_z gates at reduced fidelity, state injection avoids costly magic state distillation

factories, freeing resources for logical qubit encoding. This deliberate trade-off between performance and resource consumption allows EFT schemes to enhance algorithmic fidelity on near-term devices while circumventing the prohibitive overhead of full fault tolerance. Rather than providing a universal optimization for arbitrary circuits, state injection specifically targets applications with sparse non-Clifford gates, such as VQE and QAOA, promising candidates for demonstrating near-term quantum advantage [38–40].

However, the state injection approach remains practically limited. It still incurs additional resource overhead, including QEC resources needed to protect the ancilla logical qubit during state preparation, as well as extra quantum operations and routing required for the injection circuit [34]. Moreover, it is a general method that does not fully exploit current hardware characteristics, such as native gate fidelities, connectivity, and noise profiles [4, 5, 41, 42].

These limitations motivate the development of more efficient, hardware-informed EFT schemes, in which QEC strategies are tailored to leverage the native strengths of near-term quantum hardware. To illustrate this approach, we focus on a leading platform, trapped-ion quantum computers (TIQC) [8, 43, 44], which offer system-level features that can be directly exploited to reduce EFT overhead. In particular, we highlight three key observations: (1) TIQC’s ultra-accurate single-qubit gates already meet EFT fidelity targets (99.9999% [45]), allowing non-Clifford gates to run on bare qubits and avoiding the overhead of magic state distillation and state injection. (2) TIQC’s grid-based QCCD (quantum charge-coupled device) architecture aligns naturally with the 2D layout of surface code patches, supporting efficient logical CNOT gates. Its ion-shuttling capability enables transversal operations by bringing two surface code patches together and applying physical CNOT gates between corresponding qubits, significantly reducing the cost of multi-qubit logical operations. (3) Surface codes support dynamic encoding transitions. Existing code deformation protocols have demonstrated the possibility of transitioning between bare and logical encodings [36, 46–50]. Although these protocols are too costly for direct use, they inspire the development of runtime encoding-switch schemes compatible with on-the-fly execution.

Building on these insights, we propose iSwitch, a hybrid QEC framework for the TIQC of QCCD architecture—the next-generation device from Quantinuum—that achieves EFT with minimal resource overhead for non-Clifford gates. iSwitch employs a *temporal selective QEC* strategy. This selective protection is applied both *spatially*—encoding only qubits that require fault tolerance—and *temporally*—allowing qubits to transition between bare and logical states according to circuit requirements. Together, these strategies achieve a balanced trade-off between fidelity and resource efficiency. As shown in Fig. 1, its design consists of three synergistic modules: (1) **Runtime Encoding Protocol (Sec. 3)** enables

low-noise transitions between bare and logical qubits, allowing implementation of $R_z(\theta)$ gates without additional ancilla logical qubits. (2) **Hardware-aware Hybrid ISA (Sec. 4)** partitions the system into bare, logical, and interface regions, providing tailored instructions to coordinate operations and encoding transitions while abstracting hardware details, relieving the compiler from scheduling individual physical qubits. (3) **Compiler with Dynamic Encoding (Sec. 5)** manages encoding switches, qubit routing, and gate scheduling in a fidelity- and resource-aware way, minimizing noisy code-switch operations and optimizing space-time resources to maximize selective QEC performance.

In summary, this paper makes the following contributions:

- We propose **iSwitch**, a framework for TIQC that enables EFT by applying QEC selectively in space and time. iSwitch avoids the overhead of monolithic FTQC while satisfying EFT fidelity requirements.
- We develop a runtime protocol for in-situ encoding conversion, a hybrid ISA for managing bare–logical execution, and a compiler that co-optimizes encoding allocation and routing, together enabling the selective QEC strategy of iSwitch.
- Evaluation shows that iSwitch achieves up to 43× performance improvement across various EFT applications compared to NISQ execution, while using 50% and 33% fewer qubits than standard FTQC and previous EFT methods, respectively, under the same qubit budget, demonstrating its practical utility.

2 Background and Related Work

This section begins by providing essential background on the value of EFT applications (Sec. 2.1) and the fundamental bottlenecks of standard FT schemes (Sec. 2.2), focusing on surface codes. We then review the capabilities of TIQC hardware that make it especially suitable for near-term EFT (Sec. 2.3). Together, these elements set the foundation for our proposed solution of leveraging TIQC for EFT, which is presented in Sec. 3, 4, 5.

2.1 EFT Applications

Early fault tolerance (EFT) has emerged as the next milestone on the roadmap to full-scale FTQC, targeting near-term systems capable of executing around million robust operations with QEC protection [32]. This shift is driven by two key developments in recent years. First, the limitations of NISQ devices have become better understood [2, 3], reinforcing the view that quantum advantage is unlikely to be achieved under current noise levels without some form of fault tolerance. Second, a growing number of core FTQC building blocks have been successfully demonstrated across multiple platforms [5, 42, 51–55], suggesting that the field is now technically ready to transition from NISQ to EFT.

Despite its limited scale, EFT offers a promising path toward practical quantum advantage. Recent studies provide growing evidence that higher-accuracy VQAs, equipped with FT execution, can outperform classical heuristics by achieving better scaling and more accurate approximations to ground-state energies [56–61]. This is particularly relevant for applications in quantum chemistry [62–65] and condensed matter physics [66, 67], where classically intractable systems may become accessible through modest-scale quantum computations. Additionally, Hamiltonian simulation in areas such as quantum dynamics and quantum chemistry has demonstrated compelling accuracy improvements over classical approximation methods [62–64]. These developments underscore the practical value of EFT and motivate system-level innovations specifically designed to support this emerging computational regime.

2.2 Bottleneck of Standard FTQC Scheme

This subsection focuses on surface codes to illustrate the substantial overheads of standard FTQC protocols, which presents a major barrier to practical implementation on near-term hardware. The dominant costs arise from the space overhead of logical encoding and the space-time overhead of non-Clifford logical T gate.




Figure 2. (a) Surface code. (b) Gate Teleportation for T gate. (c) Magic state distillation.

Logical Encoding. Surface codes are a leading candidate for FTQC due to their 2D local structure, efficient decoding algorithms [49, 68, 69], and high threshold ($\sim 1\%$) [16, 70]. However, they exhibit low encoding efficiency: a square patch of $2d^2 - 1$ physical qubits encodes only a single logical qubit, where d is the code distance, and larger d is required to achieve lower logical error rates (LERs). However, they exhibit low encoding efficiency (Fig. 2(a)): a square patch of $2d^2 - 1$ physical qubits encodes only a single logical qubit, where d is the code distance, and larger d is required to achieve lower logical error rates (LERs). Simulations suggest that under a physical error rate of 10^{-3} , reaching LERs of 10^{-4} , 10^{-5} , and 10^{-6} requires distances $d = 5, 7$, and 9 , respectively [71]. In practice, however, more complex noise structures demand even larger distances. For example, Google’s experiment [4] shows that achieving an LER of $\sim 10^{-4}$ requires a distance-13 surface code, corresponding to over

300 physical qubits per logical qubit. Consequently, even a modest 30-logical-qubit program may require close to 10,000 physical qubits—well beyond the scope of near-term quantum hardware.

Logical T Gate. To support universal quantum computation, FTQC circuits are typically decomposed into the Clifford+T gate set [72]. While Clifford gates are relatively cheap to implement on surface codes [28, 71], T gates are significantly more costly. A logical T gate is applied via gate teleportation [73], using a magic state $|T\rangle = (|0\rangle + e^{i\pi/4}|1\rangle)/\sqrt{2}$. This magic state can be prepared in a surface code patch through *state injection*, which grows a physical magic state into a logical one [36, 46], but the resulting LER remains at $O(p)$, where $p = 10^{-3} \sim 10^{-4}$, insufficient for fault-tolerant use.




Figure 3. State error rate and time cost of different MSD protocols under varying qubit budgets.

To meet the required LER (e.g., 10^{-6}), the state usually undergoes magic state distillation (MSD)—a resource-intensive process that consumes multiple noisy magic states to produce one high-fidelity output. As shown in Fig. 3, We evaluated the qubit resource and distillation time cost of the $(15\text{-to-}1)_{d_x, d_z, d_m}$ MSD protocol [31]. To meet the distance-9 surface codes’s LER, this translates to approximately 2,500 physical qubits for a T gate. Furthermore, approximating a general quantum gate (e.g., $R_z(\theta)$) to precision 10^{-6} requires at least $\frac{3}{2} \log_2(10^6) \approx 30$ T gates [72], significantly increasing circuit depth and gate count.

While the above FTQC protocol is of foundational theoretical importance and has long been the prevailing approach, recent hardware advances have prompted a shift toward system-level considerations. For the less stringent requirements of EFT, much of the traditional overhead can be circumvented through system-aware design that leverages specific features of trapped-ion hardware and surface codes—as explored in the next section.

2.3 TIQC Hardware Capability

Trapped-ion quantum computers (TIQC) [8, 43] have been a leading quantum platform for over three decades [74], offering unique advantages for realizing EFT.

Ion Qubits and Gate Operations. TIQC systems encode qubits in the internal electronic states of ions confined in linear Paul traps [75], where Coulomb repulsion naturally arranges the ions into one-dimensional (1D) chains. These qubits exhibit exceptionally long coherence times—from minutes to days [43, 76]. Single-qubit (1Q) gates are implemented via tightly focused Raman lasers or global microwave fields, achieving fidelities up to 99.9999% [45], already meeting the fidelity requirements of EFT. Two-qubit (2Q) gates are realized via the Mølmer-Sørensen (MS) interaction [77, 77, 78], with fidelities reaching 99.97% [79–81]. Crucially, 2Q gates can be performed between any ion pair within a trap, enabling native all-to-all connectivity. TIQC also supports high-fidelity state preparation and measurement (SPAM) with fidelities exceeding 99.99% [82].




Figure 4. (a) QCCD architecture of TIQC, which connects 1D traps with 4-way junctions. It supports ion shuttling and swap. (b-e) The collective ion swap between (1,3) and (2,4) using ion shuttling. (f) 2Q gates (CNOT) are performed on the qubits within the same trap.

QCCD Architecture. While one-dimensional (1D) ion traps provide full connectivity within small chains, scaling them to larger systems introduces significant control complexity. To overcome this, the quantum charge-coupled device (QCCD) architecture [83–87] connects short 1D traps across a two-dimensional (2D) grid via junctions [88, 89], as illustrated in Fig. 4(a). Ions can be shuttled between 1D traps with negligible decoherence [83, 84], enabling dynamically reconfigurable connectivity over the grid. For example, Fig. 4(b-e) illustrates how the ion shuttling can collectively swap two pairs of ions (1,3) and (2,4). Once two ions are moved into a single trap, a 2Q gate (CNOT) can be performed (Fig. 2.3(f)).

Leveraging these capabilities—ultra-high-fidelity quantum operations, reconfigurable layouts, and architectural




Figure 5. Encoding transition from a bare qubit to a logical qubit under different ancilla initialization configurations. The procedure comprises one round of stabilizer measurement for gauge fixing, followed by $d - 1$ repeated rounds for error correction.

scalability—industry leaders such as Quantinuum have published clear roadmaps toward systems with hundreds of qubits [90, 91], positioning TIQC as a leading platform for near-term demonstrations of EFT.

3 Module I: Runtime Encoding Protocol

This section presents our runtime encoding strategy, inspired by gauge-fixing theory [49, 92], which converts a bare physical qubit, holding an arbitrary and program-specific quantum state, directly into a fully protected logical qubit encoded in a surface code of any target distance. Crucially, this transformation is performed *in-situ* on the data qubit itself, without the need for preparing additional ancilla logical qubits or post-selecting known ancilla states. Unlike conventional approaches such as state injection [36, 46], magic state distillation [28–30], or FT rotation-based state preparation [47], our method avoids the extra resource overhead for encoding ancillas and enables direct runtime switching of the data qubit’s encoding.

3.1 Runtime Encoding Switch

This subsection presents the runtime protocol for low-noise, in-situ conversion between bare and logical qubits using surface code patches. Fig. 5 shows the encoding switch from a bare to a logical qubit; the reverse process is symmetric and discussed at the end.

Step 1: Gauge-Aware Ancilla Initialization. At $t = 0$, we start with a bare qubit q_b in an arbitrary state and prepare

surrounding ancilla qubits for logical encoding (Fig. 5(a1, b1)). Ancilla initialization follows the geometry of the target surface code patch and the chosen logical operators X_L and Z_L . A natural choice aligns X_L and Z_L along the horizontal and vertical paths intersecting at q_b (Fig. 5(a2, b2)). Ancilla qubits along X_L are initialized in $|+\rangle$ (stabilized by X_i), and those along Z_L in $|0\rangle$ (stabilized by Z_i). Remaining ancilla qubits not on logical operators can be initialized arbitrarily in $|+\rangle$ or $|0\rangle$, providing flexibility to reduce gauge operators needing fixing and improve noise control (see Sec. 3.2).

This initialization defines $n - 1$ stabilizer generators over n qubits, leaving a single logical degree of freedom at q_b , characterized by the intersecting X and Z operators. Gauge-fixing theory ensures that this configuration preserves the bare qubit’s logical state during transfer to the surface code.

Step 2: Stabilizer Measurement and Gauge Fixing. At $t = 1$, we perform a round of stabilizer measurements on the target surface code patch, including both the region near the bare qubit and the ancilla-only area. Some stabilizers, already defined by initialization, yield deterministic $+1$ outcomes (solid operators in Fig. 5(a3, b3, c3)). The remaining stabilizers, which anti-commute with those set by initialized qubits, produce random ± 1 outcomes (dashed stabilizers in Fig. 5(a4, b4, c4)). These random outcomes do not compromise subsequent QEC and can be tracked via stabilizer frame updates or corrected with efficient gauge-fixing decoders [93, 94].

Step 3: QEC Cycles and Code Stabilization. Following gauge fixing, we perform $d - 1$ rounds of standard surface-code error correction (Fig. 5(a5, b5, c5)), where d is the target

surface code distance. These rounds further suppress noise and allow the decoder to identify and correct residual errors. At the end of this process, the original quantum state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ is promoted to a fully protected logical state $|\psi_L\rangle = \alpha|0_L\rangle + \beta|1_L\rangle$, encoded in a surface code patch of distance d . The resulting logical qubit is equivalent as one initialized via standard fault-tolerant state preparation.

Reverse Switch: Shrinking Back to Bare Qubit. Shrinking a logical qubit back to a bare qubit follows the reverse procedure. We measure and discard the ancilla qubits along X_L and Z_L (not including q_b), collapsing them in the same basis as their original initialization. For instance, qubits initialized in $|+\rangle$ are measured in the X basis (q_4, q_7 in Fig. 5(a1-a5)). Due to random measurement results in gauge fixing, measuring these ancilla qubits may introduce a global sign flip on the logical operator. For example, if the logical operator is $X_L = X_b X_4 X_7$ (Fig. 5(a1-a5)), and the product of measurement outcomes on q_4 and q_7 is -1 , the sign of the X logical operator (X_b) has flipped to -1 . In this case, we apply a correction by flipping the result of the subsequent X_b measurement or applying a Z_b gate. This ensures that the bare qubit accurately reflects the encoded logical state.

3.2 Impact of Encoding Protocols

In the encoding switch, a bare qubit can, in principle, occupy any data qubit, as the intersecting row and column naturally define valid logical X and Z operators. However, the bare qubit's placement and the initialization of surrounding ancilla qubits significantly impact the conversion's logical error rate. Hence, it is critical to identify a configuration that maximizes the fidelity of the encoding switch.

The key to suppressing errors during the switch lies in preserving as many deterministic stabilizer measurements as possible. Since error correction relies on detecting syndrome changes, stabilizers with random outcomes cannot contribute to identifying errors. For example, in Fig. 5(a3), a Z error on qubit q_4 goes undetected because the error only flips the outcome of the stabilizer $X_{1,2,4,5}$, whose measurement result is already equally likely to be ± 1 .

Moreover, whether a stabilizer is deterministic or random depends on the initialization of the ancilla qubits. For instance, in Fig. 6(c), ancilla qubits q_3, q_4, q_8 , and q_9 are initialized in the $|+\rangle$ basis. As a result, the operator $s = X_{3,4,8,9}$ stabilizes the joint state of the bare qubit and the initialized ancillas. During the gauge-fixing step shown in Fig. 6(d), all surface-code stabilizers are measured. Since these stabilizers commute with s , its syndrome is expected to remain $+1$, making s a *deterministic stabilizer*.

To maximize conversion fidelity, we introduce a new encoding protocol that initializes ancilla qubits in four triangular regions—two in the $|+\rangle$ basis and two in the $|0\rangle$ basis (see Fig. 6(c)). This arrangement ensures that all ancillas are protected by deterministic stabilizers while maintaining

balanced region sizes to provide uniform error-correction capability across the layout.

Focusing on Z errors without loss of generality, we find that only Z errors on q_8, q_{13} , and q_{18} can cause logical faults. For example, a Z_3 error flips the syndromes of both $X_{2,3}$ and $X_{3,4,8,9}$, allowing the decoder to detect and correct it unambiguously. In contrast, Z_8 and Z_9 produce identical syndrome patterns—flipping only $X_{3,4,8,9}$. Our decoder consistently attributes this pattern to Z_9 ; thus, if the actual error is Z_8 , the mis-correction results in a logical Z error during conversion.

Consequently, only a small subset of qubit errors ($Z_8, Z_{13}, Z_{18}, X_{12}, X_{13}, X_{14}$) can cause logical failure. Importantly, this implies that the probability of conversion-induced logical errors is independent of the surface code distance d , meaning that increasing the code distance does not mitigate this error. Therefore, careful hardware-aware architecture design and compiler strategies are essential to efficiently manage and minimize the overhead of encoding-switch operations.




Figure 6. Optimized runtime encoding protocol maximizing the number of deterministic stabilizers. (a)(b) Bare qubit placement defining logical operators. (c) Ancilla initialization pattern with $|+\rangle$ and $|0\rangle$ maximizes deterministic stabilizers. (d) Resulting stabilizer measurement outcomes.

4 Module II: Hybrid Quantum ISA

To enable dynamic switching between bare and protected execution, we introduce a hybrid quantum instruction set architecture (ISA) and corresponding microarchitectural designs tailored to TIQC. As illustrated in Fig. 1, the hardware is partitioned into three regions: a *bare qubit region*, a *logical qubit region*, and a *boundary region*. The boundary region serves as a reconfigurable interface between bare and logical domains, where each patch in this region can be interpreted

Table 1. Hybrid ISA for TIQC: Bare, Logical, and Encoding Switch Instructions

Category	Instruction	Semantics and Purpose
Bare Qubit Instructions	BareMove_Vertical	Move a bare ion qubit up/down across rows of trap junctions.
	BareMove_Horizontal	Move a bare ion qubit left/right across columns of trap junctions.
	Bare1Q_Gate	High-fidelity native single-qubit gates executed on bare qubits.
Logical Qubit Instructions	LogicMove_Vertical	Move a logical patch up/down one step along the trap QEC patch grid.
	LogicMove_Horizontal	Move a logical patch left/right one step along the trap QEC patch grid.
	LogicCX_Transversal	Perform a logical CNOT gate transversally between two logical on the two overlapped patches.
Bare-Logical Instructions	CodeSwitch_B2L	Convert a bare qubit into a logical patch by growing stabilizers region.
	CodeSwitch_L2B	Convert a logical patch into a bare qubit by shrinking its stabilizer region.

either as a logical surface code block or as a bare qubit before enlarging or after shrinking.

4.1 Hybrid ISA Summary

Our hybrid ISA defines three distinct categories of quantum instructions, corresponding to three execution domains: *bare qubit operations*, *logical patch operations*, and *bare-logical switching operations*. This division aligns with our hardware layout and supports efficient compilation and scheduling. A summary of the hybrid ISA is provided in Table 1.

Bare Qubit Instructions. These instructions operate on unencoded qubits in the bare region. Designed to leverage the native strengths of TIQC, they enable high-fidelity operations with minimal overhead. BareMove perform ion swaps across traps to dynamically reconfigure bare qubit layouts. Bare1Q_Gate allows arbitrary native rotations to be directly applied to bare qubits, avoiding the need for QEC overhead when protection is not required.

Logical Patch Instructions. These instructions operate on encoded qubits within surface code patches. To enable flexible layout and interactions, logical qubits must be movable while preserving their stabilizer structure. The LogicMove instructions allow movement across the 2D QCCD grid, supporting logical gate reconfiguration without requiring teleportation or decoding.

Bare-Logical Switching Instructions. To support hybrid execution, the ISA includes two dedicated instructions — CodeSwitch_B2L and CodeSwitch_L2B — for runtime conversion between encoding domains. These operations are restricted to the boundary region, which simplifies scheduling and localizes the coordination between bare and logical regions. This constraint ensures modularity and minimizes inter-domain coupling.

Design Rationale. The ISA is deliberately minimal and modular, reducing architectural complexity while maintaining flexibility. By isolating switching operations to explicit boundary instructions, the compiler can optimize independently within each domain, ensuring scalable and error-aware scheduling. The dual-granularity design, which uses

bare qubits for high-fidelity, low-overhead operations, and logical qubits for protected execution, enables tunable trade-offs between fidelity, resource usage, and performance, laying a robust foundation for dynamic compilation (Sec. 5).

4.2 Microarchitecture Designs

Our architectural design targets 2D junction-based ion trap systems (Sec. 2.3), inspired by recent Quantinuum hardware and surface-code-based ion trap schemes proposed in [88, 89, 95]. We divide the hardware into bare and logical regions, and ensure hybrid ISA can be implemented efficiently across both domains. We divide the hardware into bare and logical regions, and ensure instructions in Table 1 can be implemented efficiently across both domains.




Figure 7. (a) Mapping surface-code patches onto a 2D QCCD grid. (b) Patches are moved to clear routing channels. (c) Patches are aligned into an overlapping region for interaction. (d) Return the patch to perform transversal CNOT.

LogicMove. Fig. 7(a) illustrates the implementation of the surface code in our architecture, while Fig. 7(b–c) shows the fine-grained steps involved in logical patch movement. The highlighted yellow regions denote computation zones that support either two-qubit physical gates or single-qubit ion measurements. Each one-dimensional (1D) trap can host multiple physical qubits; for clarity of exposition, we initially

place a single data qubit in each 1D trap. Ancilla ions are placed along vertical junctions and serve as syndrome qubits for stabilizer measurements. Movement of adjacent logical patches is realized via ion shuttling across trap junctions.

LogicCX_Transversal. To implement transversal CNOT gates between two logical patches, we leverage 1D traps that can contain multiple ions. As shown in Fig. 7(d), logical patches are moved into proximity such that each 1D trap hosts a pair of qubits—one from each patch. CNOT gates are then applied in parallel across all pairs. After execution, the patches are separated by reversing the movement steps. This approach ensures FT logical CNOT without requiring teleportation or measurement-based gates.

CodeSwitch. At the boundary region, `CodeSwitch_B2L` and `CodeSwitch_L2B` instructions are invoked to transition data between unencoded and protected forms (Fig. 8). These operations restrict deformation to boundary region, which significantly simplifies both architectural design and scheduling. As a result, intra-region operations can remain fully independent, avoiding compiler-level entanglement between instruction types.




Figure 8. Switch between bare (P_1) and logical qubits ($Q_{2,3}$).

5 Module III: Compiler Support

Building on the hybrid ISA and region-aware architecture introduced in Sec. 4, we develop a fidelity- and resource-aware compiler that manages both the encoding state and spatial placement of program qubits. It addresses two main challenges: (1) *dynamic encoding allocation*, deciding when and where qubits reside in logical versus bare form, and (2) *routing and co-location of logical patches* to enable two-qubit gates efficiently.

When the logical patch region can encode all program qubits, the compiler can perform `CodeSwitch` in-place, removing the need for a separate physical qubit region. "In-place" means the data qubit does not need to move to the boundary. When ancilla resources are sufficient, the ancillas

released after an L2B operation for logical qubit Q_1 can remain locally associated with Q_1 and be reused in its next B2L, instead of being reassigned to another qubit Q_2 . A qubit can temporarily switch to bare form for a single-qubit R_z gate and immediately switch back to a logical patch for protection, avoiding any physical movement or routing overhead.

In practice, however, the number of active surface code patches is limited, constraining the "logical register" space, while the bare region provides flexible, lower-fidelity memory. Additionally, grid-based adjacency restricts gate execution, often requiring explicit qubit movement.

To handle these constraints, the compiler leverages abstractions inspired by classical register allocation and NISQ-era routing. The resulting compilation passes generate low-overhead, error-aware schedules that respect encoding constraints while balancing fidelity, locality, and resource usage.

5.1 Bare-Logical Encoding Allocation

A central challenge in compiling programs for iSwitch is dynamically deciding which qubits should reside in logical patches and which can remain as bare qubits, analogous to classical register allocation with limited fast-access registers.

In iSwitch, logical patches function as a finite register file, while bare qubits serve as main memory. Program operations fall into two categories: (1) high-fidelity single-qubit non-Clifford gates on bare qubits, and (2) two-qubit gates requiring both qubits in logical form. Minimizing *spills*—runtime conversions that incur overhead and introduce additional error—is the primary compilation goal. We model encoding allocation as a temporal coloring problem over the program's qubit usage graph and employ a greedy linear-scan algorithm inspired by classical register allocation.

At each timestep, a two-qubit gate triggers logical encoding for both operands; if capacity is insufficient, idle logical qubits are shrunk to bare form to free space. Single-qubit gates are executed on bare qubits, postponing conversion until necessary. This temporal reuse strategy reduces conversions and maximizes fidelity.

The compiler maintains a pool of logical qubits limited by the available logical region. Conversions are applied only when required, and logical qubits are proactively shrunk once protection is no longer needed, freeing space for subsequent allocations. This mirrors classical register allocation and spilling, while incorporating a fidelity-aware model in which each conversion introduces additional noise.

5.2 Routing: Connectivity Management

Beyond encoding allocation, the compiler must also ensure that logical qubits involved in 2Q operations are physically adjacent on the ion trap grid. To this end, we adopt a SABRE-inspired routing strategy [96], treating logical patches as movable units and scheduling swaps to co-locate interacting qubits. Given the 2D layout of the logical patch region, we




Figure 9. (a) Experiment results on a Quantinuum’s H1 device to measure noise channel parameters. (b)–(d) Logical error rate simulations using the experimentally calibrated noise model, showing the logical error rates for (b) transversal 2-qubit gates, (c) logical patch movement, and (d) code conversion operations across varying code distances.

construct a logical coupling graph and apply SABRE to compute low-cost swap schedules. Each proposed Swap operation is translated into our hybrid ISA using `LogicMove_Vertical` or `LogicMove_Horizontal` instructions, depending on the alignment of the operand patches.

A similar routing strategy is applied within the bare region, where qubit movement is implemented via physical ion swaps. At the interface between bare and logical zones, we incorporate conversion constraints into the routing logic to jointly optimize conversion and movement steps. This co-optimization enables efficient spatial scheduling while maintaining correctness and minimizing decoherence.

6 Evaluation

6.1 Experiment Setup

Baselines. We compare iSwitch against three baselines:

- (1) **NISQ-Bare:** A NISQ-only baseline in which all operations are executed on bare physical qubits without quantum error correction. This approach incurs minimal qubit overhead but is highly vulnerable to noise.
- (2) **MSD-Logical:** A fully fault-tolerant baseline that uses surface code with distance- d encoding for all program qubits. All gates are implemented fault-tolerantly. Non-Clifford $R_Z(\theta)$ gates are decomposed into T gates using Gridsynth [72], and each T gate is executed via a distilled magic state prepared using the $(15\text{-to-}1)_{13,5,5}$ protocol [31]. As a result, MSD-Logical must allocate a significant portion of its qubit budget to magic state distillation (MSD), thereby reducing the number of qubits available for encoding program logical qubits.
- (3) **INJ-Logical:** A partially fault-tolerant baseline that implements non-Clifford $R_Z(\theta)$ gates via state injection [34, 36]. The protocol first prepares the $R_Z(\theta)$ state on a small $d = 3$ code patch using post-selection, then enlarges the patch to form an ancilla logical qubit compatible with the program qubits. The logical $R_Z(\theta)$ gate is then applied via teleportation using logical CNOT gates and measurements.

Noise Model. To simulate realistic trapped-ion quantum computing (TIQC) behavior, we use a gate-level noise model calibrated with recent experimental benchmarks [8, 44, 45,

97]. The error rates are derived from our randomized benchmarking experiments on the Quantinuum’s H1 system, as shown in Fig. 9(a), with single-qubit gates at 1×10^{-5} , two-qubit gates at 10^{-3} , SPAM (state preparation and measurement) errors at 10^{-3} , idling errors at 10^{-5} and shuttling errors at 10^{-4} , consistent with current TIQC performance [98, 99].

Benchmarks. We evaluate our approach on variational quantum algorithms using the *unitary coupled-cluster singles and doubles* (UCCSD) ansatz [100], a standard choice for VQAs. In particular, our benchmark design follows recent EFT studies [34], ensuring comparability with state-injection-based evaluation. We focus on ground-state energy estimation tasks from the physics models and quantum chemistry [66, 67], relevant to EFT quantum computing.

Metrics. We adopt two primary metrics for end-to-end performance evaluation: (1) *VQA energy* E_s , which denotes the final converged energy produced by the VQA circuit. As proposed in [101], lower energy values correspond to better performance. To quantify improvement, we compute the ratio of the energy gap between the baseline and the ideal (noiseless) energy to that achieved by iSwitch. A higher ratio indicates greater performance gain. (2) *Program fidelity* f , defined as the probability that a benchmark completes without any logical error. Higher fidelity implies better reliability and correctness. To quantify the improvement of iSwitch over baseline A, we define the fidelity improvement ratio as $\gamma = \frac{1-f_A}{1-f_{iSwitch}}$. A larger value of γ indicates that iSwitch achieves a greater reduction of infidelity.

Evaluation Setting. To balance simulation accuracy and efficiency, we follow simulation strategies from previous works [34, 36] for small- and large-scale circuits. Evaluation was conducted on a server with an AMD EPYC 9534 64-core processor and 1.48 TB of memory.

Small-scale circuits (≤ 10 qubits). For small circuits, we use the Qiskit-Aer simulator [102] in density matrix mode to model noise based on realistic hardware parameters.

Large-scale circuits (> 10 qubits). Discrete parameter optimization is performed with a classical genetic algorithm [103, 104]




Figure 10. VQA energy comparison of iSwitch ($d = 9$) and NISQ-Bare. Ideal denotes the noiseless VQA circuit.

using tournament selection, crossover, and mutation over 50 generations with a population size of 64.

6.2 Simulation of Logical Operations

We first validate our noise model on the Quantinuum H1 device [98], and then use the validated model to simulate the LERs of the operations (LogicCX, LogicMove, CodeSwitch) defined in Table 1 across various code distances.

1. Noise model validation. In Fig. 9(a), we compare the overall error rate accumulation as a function of operation count for different physical qubit operations, including CNOT errors, SPAM errors, idling errors (with and without qubit shuttling), and single-qubit gate errors. The close agreement between experimental *markers* and simulation *lines* confirms the accuracy of our model, providing a solid foundation for subsequent logical operation simulations.

2. Logical error rate simulation. In Fig. 9(b–d) We simulate the LERs of three logical operations with Stim [105] and Pymatching [69]. Our key observations are as follows:

- (i) The LERs of LogicCX and LogicMove decrease with increasing code distance, demonstrating fault-tolerance. In contrast, CodeSwitch remains independent of code distance, reflecting the partial fault-tolerance inherent in our iSwitch.
- (ii) In Fig. 9(c), the LER of patch movement increases with the number of rows or columns traversed. Here, "move 1" corresponds to moving a patch to overlap with an adjacent patch, while "move 0" indicates that the patch remains idle. This trend arises from the accumulation of idling errors in the absence of intermediate QEC. However, a single distance-2 move still introduces fewer errors than two consecutive distance-1 moves, which can be attributed to the extra patch movement required to clear the path and return (Fig. 7(b)(d)).
- (iii) *Comparison of runtime encoding strategies (Fig. 9d).* We consider three protocols: (1) *iSwitch*, (2) *Corner Enlargement*,

which performs the switch at the patch corner, and (3) *Random Initial State*, which targets the patch center but randomly initializes ancilla qubits in $|0\rangle$ or $|+\rangle$ states during enlargement (see Sec. 3.1). Our in-center-switch strategy achieves the lowest LER by maximizing the number of deterministic stabilizers around the injected state, thereby protecting the conversion process. In contrast, random initialization performs progressively worse for larger codes, as the absence of deterministic stabilizers allows nearly any qubit error to propagate into a logical error, and the probability of such errors increases with code size.

6.3 Comparison to NISQ Approach

We evaluate the improvement of iSwitch, which uses a surface code with $d = 9$, over NISQ-Bare by comparing the final converged VQA energies, as shown in Fig. 10, which reflect end-to-end performance. An "Ideal" reference line represents the noiseless circuit performance. To quantify improvement, we define the energy improvement ratio as the gap between NISQ-Bare and the ideal baseline, divided by the gap under iSwitch. iSwitch consistently achieves lower VQA energies. For small-scale cases (Fig. 10(a)–(d)), the improvement ratios are $6.89\times$, $13.1\times$, $4.64\times$, and $7.38\times$, respectively. For large-scale benchmarks (Fig. 10(e)–(h)), the gains increase to $4.34\times$, $43.4\times$, $10.8\times$, and $23.6\times$.

This improvement arises from three key factors. First, in iSwitch, all 2Q gates are QEC-protected with LERs below 10^{-6} , whereas in NISQ-Bare, they remain unprotected with LERs near 10^{-3} . This difference in fidelity significantly reduces accumulated noise. Second, in VQA circuits, the number of 2Q gates is much larger than the number of non-Clifford gates, making them the primary source of error. As the circuit size increases, the gap between the number of 2Q gates and non-Clifford gates grows, leading to a greater




Figure 11. Program fidelity comparison between iSwitch (orange) and other QEC approaches, including state injection (purple) and magic state distillation (red), across varying device sizes and benchmarks.




Figure 12. Ablation study of iSwitch. (a) Impact of varying code distance on program fidelity. (b) Impact of omitting Clifford+T decomposition on circuit depth reduction. (c) Impact of varying encoding ratio on the number of encoding conversions.

reduction in noise for larger circuits in iSwitch. Third, our compiler minimizes noisy transitions between bare and logical qubits, trading a small number of code-switching operations to protect all 2Q gates on logical qubits, thus avoiding noisy 2Q operations on bare qubits.

6.4 Comparison to QEC Approaches

In this section, we evaluate the improvement of iSwitch over two current QEC methods: fully-QEC MSD-Logical and partial-QEC INJ-Logical. In Fig. 11, we compare the program fidelity of iSwitch (orange), INJ-Logical (purple), and MSD-Logical (red) across four benchmarks.

Comparison to MSD-Based Method. As shown in Fig. 11, MSD-Logical requires, on average, $2.06\times$ more qubits to achieve the same fidelity as iSwitch. While MSD-Logical can achieve higher fidelity with sufficient qubit resources (as it is a fully fault-tolerant method), iSwitch remains superior for larger programs ($n > 30$) with equivalent qubit resources.

This improvement stems from two main factors: First, MSD-Logical allocates a significant portion of its qubit budget to magic state factories, forcing the use of lower-distance encoding for program qubits. This results in higher LER of 2Q gates due to insufficient protection. For an example, with a budget of 4860 qubits, MSD-Logical uses 2594 qubits to construct the T-factory, leaving only enough qubits to encode program qubits using distance-6 surface codes.

Second, MSD-Logical requires decomposing all non-Clifford single-qubit gates (e.g., $R_Z(\theta)$) into Clifford+T gates, leading to significant circuit depth overhead. The increased depth exacerbates noise accumulation, as logical qubits must idle while a target qubit sequentially executes multiple T gates.

Comparison to Injection-Based Method. As shown in Fig. 11, INJ-Logical requires, on average, $1.49\times$ more qubits to achieve the same fidelity as iSwitch. Both methods reach a fidelity ceiling, as they are partial fault-tolerant, and overall program fidelity is constrained by the non-FT R_Z gates. The advantages of iSwitch arise from the following factors:

First, INJ-Logical requires all program qubits to be fully encoded and additionally relies on approximately 50% more logical qubits [34] for ancilla qubits to prepare and store injected $R_Z(\theta)$ states. This approach reduces flexibility of encoding ratio compared to iSwitch’s hybrid design and requires extra resources for implementing the R_Z gate.

Second, injection-based implementations of $R_Z(\theta)$ often require multiple ancillary states, such as $R_Z(\theta)$, $R_Z(2\theta)$, and $R_Z(4\theta)$, to probabilistically construct a single $R_Z(\theta)$ rotation. The state injection protocol itself is probabilistic, and there is a 50% chance that an injected $R_Z(\theta)$ state will result in $R_Z(-\theta)$, requiring the injection of an additional $R_Z(2\theta)$ state. This probabilistic nature means that, on average, two injected R_Z states are needed to construct a single $R_Z(\theta)$ gate, with the number of states to prepare being uncertain in advance.

Third, the success probability of these state injection protocols decreases as the surface code distance increases. As a result, injection methods often initialize the state on a small code (e.g., $d = 3$) and gradually enlarge it, which increases circuit depth and exposure to noise. Additionally, because post-selection is used in preparing the $R_Z(\theta)$ states, the probability of successfully generating the state decreases with increasing code distance. This gradual enlargement of the surface code increases circuit depth, and during this process, the generated R_Z state is more susceptible to noise due to insufficient code distance.

Scalability analysis. We observe an empirical *diminishing-return point* beyond which the fully fault-tolerant (MSD-Logical) execution begins to match or exceed the performance of iSwitch, such that allocating additional physical qubits no longer yields meaningful fidelity gains for EFT. For the evaluated 20 logical-qubit workloads, this transition occurs around $\sim 5k$ physical qubits (Fig. 11 (a)(c)).

Importantly, this transition is *not* a fixed qubit-scale threshold, but shifts with the *program scale*. For 30 logical-qubit workloads, the diminishing point moves to around $\sim 7k$ physical qubits. The performance gap is governed by the effective physical-qubit budget per logical qubit. For a fixed device size, increasing the number of logical qubits reduces this per-logical budget and therefore delays the regime in which Full-FT overtakes EFT.

6.5 Component Justification Analysis.

Impact of Surface Code Distance. We study how varying the surface code distance affects the performance of iSwitch, as measured by the fidelity improvement ratio over NISQ-Bare (Fig. 12(a)). (1) The results indicate that increasing the code distance beyond $d = 9$ yields only marginal improvements in program fidelity. This is because Beyond $d = 9$, the logical CNOT error becomes sufficiently suppressed, and conversion fidelity becomes the dominant performance bottleneck, which further increasing the code distance cannot mitigate. (2) In contrast, small-distance surface codes fail to provide adequate protection. At $d = 5$, the fidelity improvement is minimal, and at $d = 3$, performance is even worse than the NISQ-Bare baseline—resulting in higher VQA energy—due to the overhead of syndrome measurements and conversion-related noise.

Impact of omitting Clifford+T Decomposition. In iSwitch, non-Clifford single-qubit gates (e.g., $R_Z(\theta)$) are executed directly on bare qubits, avoiding Clifford+T decomposition. This significantly reduces circuit depth: as shown in Fig. 12(b), conventional MSD-based methods incur $14.9\times$ – $27.8\times$ depth increases depending on decomposition accuracy ϵ . The shorter circuits lower qubit idling errors, reducing noise accumulation from waiting for T-state distillation. Furthermore, because each $R_Z(\theta)$ gate is decomposed into multiple T gates, maintaining low accumulated error requires high-fidelity T

states. This, in turn, demands a high-quality MSD factory, which can consume substantial qubit resources.

Impact of Logical Encoding Ratio. We evaluate how limiting the number of simultaneously active logical qubits affects performance. Specifically, we vary the *logical encoding ratio*—defined as the maximum number of concurrently encoded logical qubits divided by the total number of program qubits—and measure its impact on runtime behavior. As shown in Fig. 12(c), reducing the logical encoding ratio increases the number of encoding-switch operations required. This is expected: with fewer logical qubits allowed, qubits must more frequently transition between bare and logical forms to satisfy gate dependencies, increasing swap and routing activity across architectural regions.

7 Conclusion

We introduced iSwitch, a compiler-assisted selective QEC framework for TIQC. By in-situ encoding switching, it executes single-qubit non-Clifford gates with zero extra ancilla logical overhead and maintains high-fidelity two-qubit operations, reducing qubit and conversion cost compared to fully fault-tolerant and existing early fault-tolerant methods. These results demonstrate practical, near-term fault tolerance on hardware-constrained trapped-ion devices.

8 Acknowledgments

We thank the anonymous reviewers for their constructive feedback, and Ang Li from PNNL for helpful discussions. This material is based upon work supported by the U.S. Department of Energy, Office of Science, National Quantum Information Science Research Centers, Quantum Science Center (QSC). This work was supported in part by NSF Grants 2048144, NSF 2422169, NSF 2427109, NSF 2519029 and PHY-2110421, as well as ARO Grant W911NF-20-1-0037. Additional support is provided by NSF National Quantum Virtual Laboratory (NQVL) under NSF 2435382 and the NSF Challenge Institute for Quantum Computation (CIQC) under NSF 2016245. This work was also supported in part by the Cisco Research Fund and benefited from NVIDIA CUDA-Q and Quantinuum Systems hardware promotional offers.

References

- [1] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell *et al.*, “Quantum supremacy using a programmable superconducting processor,” *Nature*, vol. 574, no. 7779, pp. 505–510, 2019.
- [2] D. Stilck França and R. Garcia-Patron, “Limitations of optimization algorithms on noisy quantum devices,” *Nature Physics*, vol. 17, no. 11, pp. 1221–1227, 2021.
- [3] G. De Palma, M. Marvian, C. Rouzé, and D. S. França, “Limitations of variational quantum algorithms: a quantum optimal transport approach,” *PRX Quantum*, vol. 4, no. 1, p. 010309, 2023.
- [4] “Suppressing quantum errors by scaling a surface code logical qubit,” *Nature*, vol. 614, no. 7949, pp. 676–681, 2023.

- [5] R. Acharya, L. Aghababaie-Beni, I. Aleiner, T. I. Andersen, M. Ansmann, F. Arute, K. Arya, A. Asfaw, N. Astrakhantsev, J. Atalaya, R. Babbush, D. Bacon, B. Ballard, J. C. Bardin, J. Bausch, A. Bengtsson, A. Bिल्mes, S. Blackwell, S. Boixo, G. Bortoli, A. Bourassa, J. Bovaird, L. Brill, M. Broughton, D. A. Browne, B. Buchea, B. B. Buckley, D. A. Buell, T. Burger, B. Burkett, N. Bushnell, A. Cabrera, J. Campero, H.-S. Chang, Y. Chen, Z. Chen, B. Chiaro, D. Chik, C. Chou, J. Claes, A. Y. Cleland, J. Cogan, R. Collins, P. Conner, W. Courtney, A. L. Crook, B. Curtin, S. Das, A. Davies, L. D. Lorenzo, D. M. Debroy, S. Demura, M. Devoret, A. D. Paolo, P. Donohoe, I. Drozdov, A. Dunsworth, C. Earle, T. Edlich, A. Eickbusch, A. M. Elbag, M. Elzouka, C. Erickson, L. Faoro, E. Farhi, V. S. Ferreira, L. F. Burgos, E. Forati, A. G. Fowler, B. Foxen, S. Ganjam, G. Garcia, R. Gasca, Élie Genois, W. Giang, C. Gidney, D. Gilboa, R. Gosula, A. G. Dau, D. Graumann, A. Greene, J. A. Gross, S. Habegger, J. Hall, M. C. Hamilton, M. Hansen, M. P. Harrigan, S. D. Harrington, F. J. H. Heras, S. Heslin, P. Heu, O. Higgott, G. Hill, J. Hilton, G. Holland, S. Hong, H.-Y. Huang, A. Huff, W. J. Huggins, L. B. Joffe, S. V. Isakov, J. Iveland, E. Jeffrey, Z. Jiang, C. Jones, S. Jordan, C. Joshi, P. Juhas, D. Kafri, H. Kang, A. H. Karamlou, K. Kechedzhi, J. Kelly, T. Khair, T. Khattar, M. Khezri, S. Kim, P. V. Klimov, A. R. Klots, B. Koblin, P. Kohli, A. N. Korotkov, F. Kostritsa, R. Kothari, B. Kozlovskii, J. M. Kreikebaum, V. D. Kurilovich, N. Lacroix, D. Landhuis, T. Lange-Dei, B. W. Langley, P. Laptev, K.-M. Lau, L. L. Guevel, J. Ledford, K. Lee, Y. D. Lensky, S. Leon, B. J. Lester, W. Y. Li, Y. Li, A. T. Lill, W. Liu, W. P. Livingston, A. Lochar, E. Lucero, D. Lundahl, A. Lunt, S. Madhuk, F. D. Malone, A. Maloney, S. Mandrá, L. S. Martin, S. Martin, O. Martin, C. Maxfield, J. R. McClean, M. McEwen, S. Meeks, A. Megrant, X. Mi, K. C. Miao, A. Mieszala, R. Molavi, S. Molina, S. Montazeri, A. Morvan, R. Movassagh, W. Mruczkiewicz, O. Naaman, M. Neeley, C. Neill, A. Nersisyan, H. Neven, M. Newman, J. H. Ng, A. Nguyen, M. Nguyen, C.-H. Ni, T. E. O'Brien, W. D. Oliver, A. Opremcak, K. Ottosson, A. Petukhov, A. Pizzuto, J. Platt, R. Potter, O. Pritchard, L. P. Pryadko, C. Quintana, G. Ramachandran, M. J. Reagor, D. M. Rhodes, G. Roberts, E. Rosenberg, E. Rosenfeld, P. Roushan, N. C. Rubin, N. Saei, D. Sank, K. Sankaragomathi, K. J. Satzinger, H. F. Schurkus, C. Schuster, A. W. Senior, M. J. Shearn, A. Shorter, N. Shutty, V. Shvarts, S. Singh, V. Sivak, J. Skrzynski, S. Small, V. Smelyanskiy, W. C. Smith, R. D. Somma, S. Springer, G. Sterling, D. Strain, J. Suchard, A. Szasz, A. Sztein, D. Thor, A. Torres, M. M. Torunbalci, A. Vaishnav, J. Vargas, S. Vdovichev, G. Vidal, B. Villalonga, C. V. Heidweiller, S. Waltman, S. X. Wang, B. Ware, K. Weber, T. White, K. Wong, B. W. K. Woo, C. Xing, Z. J. Yao, P. Yeh, B. Ying, J. Yoo, N. Yosri, G. Young, A. Zalcman, Y. Zhang, N. Zhu, and N. Zorbrist, "Quantum error correction below the surface code threshold," 2024.
- [6] H.-L. Huang, D. Wu, D. Fan, and X. Zhu, "Superconducting quantum computing: a review," *Science China Information Sciences*, vol. 63, pp. 1–32, 2020.
- [7] S. Bravyi, O. Dial, J. M. Gambetta, D. Gil, and Z. Nazario, "The future of quantum computing with superconducting qubits," *Journal of Applied Physics*, vol. 132, no. 16, 2022.
- [8] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, "Trapped-ion quantum computing: Progress and challenges," *Applied physics reviews*, vol. 6, no. 2, 2019.
- [9] J.-S. Chen, E. Nielsen, M. Ebert, V. Inlek, K. Wright, V. Chaplin, A. Maksymov, E. Páez, A. Poudel, P. Maunz *et al.*, "Benchmarking a trapped-ion quantum computer with 29 algorithmic qubits," *arXiv preprint arXiv:2308.05071*, 2023.
- [10] J. Wurtz, A. Bylinskii, B. Braverman, J. Amato-Grill, S. H. Cantu, F. Huber, A. Lukin, F. Liu, P. Weinberg, J. Long *et al.*, "Aquila: Quera's 256-qubit neutral-atom quantum computer," *arXiv preprint arXiv:2306.11727*, 2023.
- [11] K. Temme, S. Bravyi, and J. M. Gambetta, "Error mitigation for short-depth quantum circuits," *Physical review letters*, vol. 119, no. 18, p. 180509, 2017.
- [12] S. Endo, S. C. Benjamin, and Y. Li, "Practical quantum error mitigation for near-future applications," *Physical Review X*, vol. 8, no. 3, p. 031027, 2018.
- [13] A. R. Calderbank and P. W. Shor, "Good quantum error-correcting codes exist," *Physical Review A*, vol. 54, no. 2, p. 1098, 1996.
- [14] H. Bombin and M. A. Martin-Delgado, "Topological quantum distillation," *Physical review letters*, vol. 97, no. 18, p. 180501, 2006.
- [15] A. Y. Kitaev, "Fault-tolerant quantum computation by anyons," *Annals of physics*, vol. 303, no. 1, pp. 2–30, 2003.
- [16] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, "Surface codes: Towards practical large-scale quantum computation," *Physical Review A*, vol. 86, no. 3, p. 032324, 2012.
- [17] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, "High-threshold and low-overhead fault-tolerant quantum memory," *Nature*, vol. 627, no. 8005, pp. 778–782, 2024.
- [18] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information*. Cambridge university press, 2010.
- [19] D. Gottesman, "Theory of fault-tolerant quantum computation," *Physical Review A*, vol. 57, no. 1, p. 127, 1998.
- [20] P. W. Shor, "Fault-tolerant quantum computation," in *Proceedings of 37th conference on foundations of computer science*. IEEE, 1996, pp. 56–65.
- [21] N. Koukoulekidis, S. Wang, T. O'Leary, D. Bultrini, L. Cincio, and P. Czarnik, "A framework of partial error correction for intermediate-scale quantum computers," *arXiv preprint arXiv:2306.15531*, 2023.
- [22] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM review*, vol. 41, no. 2, pp. 303–332, 1999.
- [23] L. K. Grover, "A fast quantum mechanical algorithm for database search," in *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, 1996, pp. 212–219.
- [24] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. Sawaya *et al.*, "Quantum chemistry in the age of quantum computing," *Chemical reviews*, vol. 119, no. 19, pp. 10 856–10 915, 2019.
- [25] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller, "Practical quantum advantage in quantum simulation," *Nature*, vol. 607, no. 7920, pp. 667–676, 2022.
- [26] C. Gidney and M. Ekerå, "How to factor 2048 bit rsa integers in 8 hours using 20 million noisy qubits," *Quantum*, vol. 5, p. 433, 2021.
- [27] B. Eastin and E. Knill, "Restrictions on transversal encoded quantum gate sets," *Physical review letters*, vol. 102, no. 11, p. 110502, 2009.
- [28] S. Bravyi and A. Kitaev, "Universal quantum computation with ideal clifford gates and noisy ancillas," *Physical Review A—Atomic, Molecular, and Optical Physics*, vol. 71, no. 2, p. 022316, 2005.
- [29] S. Bravyi and J. Haah, "Magic-state distillation with low overhead," *Physical Review A—Atomic, Molecular, and Optical Physics*, vol. 86, no. 5, p. 052329, 2012.
- [30] J. Haah and M. B. Hastings, "Codes and protocols for distilling t , controlled- s , and toffoli gates," *Quantum*, vol. 2, p. 71, 2018.
- [31] D. Litinski, "Magic state distillation: Not as costly as you think," *Quantum*, vol. 3, p. 205, 2019.
- [32] J. Preskill, "Beyond nisq: The megaquop machine," 2025.
- [33] —, "Quantum computing in the nisq era and beyond," *Quantum*, vol. 2, p. 79, 2018.
- [34] S. Dangwal, S. Vittal, L. M. Seifert, F. T. Chong, and G. S. Ravi, "Variational quantum algorithms in the era of early fault tolerance," *arXiv preprint arXiv:2503.20963*, 2025.
- [35] A. Katabarwa, K. Gratsea, A. Caesura, and P. D. Johnson, "Early fault-tolerant quantum computing," *PRX quantum*, vol. 5, no. 2, p. 020101, 2024.
- [36] L. Lao and B. Criger, "Magic state injection on the rotated surface code," in *Proceedings of the 19th ACM International Conference on Computing Frontiers*, 2022, pp. 113–120.

- [37] R. Toshio, Y. Akahoshi, J. Fujisaki, H. Oshima, S. Sato, and K. Fujii, “Practical quantum advantage on partially fault-tolerant quantum computer,” *Physical Review X*, vol. 15, no. 2, p. 021057, 2025.
- [38] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio *et al.*, “Variational quantum algorithms,” *Nature Reviews Physics*, vol. 3, no. 9, pp. 625–644, 2021.
- [39] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, “A variational eigenvalue solver on a photonic quantum processor,” *Nature communications*, vol. 5, no. 1, p. 4213, 2014.
- [40] E. Farhi, J. Goldstone, and S. Gutmann, “A quantum approximate optimization algorithm,” *arXiv preprint arXiv:1411.4028*, 2014.
- [41] Y. Zhao, Y. Ye, H.-L. Huang, Y. Zhang, D. Wu, H. Guan, Q. Zhu, Z. Wei, T. He, S. Cao, F. Chen, T.-H. Chung, H. Deng, D. Fan, M. Gong, C. Guo, S. Guo, L. Han, N. Li, S. Li, Y. Li, F. Liang, J. Lin, H. Qian, H. Rong, H. Su, L. Sun, S. Wang, Y. Wu, Y. Xu, C. Ying, J. Yu, C. Zha, K. Zhang, Y.-H. Huo, C.-Y. Lu, C.-Z. Peng, X. Zhu, and J.-W. Pan, “Realization of an error-correcting surface code with superconducting qubits,” *Physical Review Letters*, vol. 129, no. 3, p. 030501, 2022.
- [42] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter, J. P. Bonilla Ataides, N. Maskara, I. Cong, X. Gao, P. Sales Rodriguez, T. Karolyshyn, G. Semeghini, M. J. Gullans, M. Greiner, V. Vuletić, and M. D. Lukin, “Logical quantum processor based on reconfigurable atom arrays,” *Nature*, vol. 626, no. 7997, pp. 58–65, 2024.
- [43] T. Strohm, K. Wintersperger, F. Dommert, D. Basilewitsch, G. Reuber, A. Houshanov, T. Ehmer, D. Vodola, and S. Luber, “Ion-based quantum computing hardware: Performance and end-user perspective,” *arXiv preprint arXiv:2405.11450*, 2024.
- [44] S. A. Moses, C. H. Baldwin, M. S. Allman, R. Ancona, L. Ascarrunz, C. Barnes, J. Bartolotta, B. Bjork, P. Blanchard, M. Bohn *et al.*, “A race-track trapped-ion quantum processor,” *Physical Review X*, vol. 13, no. 4, p. 041052, 2023.
- [45] T. Harty, D. Allcock, C. J. Ballance, L. Guidoni, H. Janacek, N. Linke, D. Stacey, and D. Lucas, “High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit,” *Physical review letters*, vol. 113, no. 22, p. 220501, 2014.
- [46] Y. Li, “A magic state’s fidelity can be superior to the operations that created it,” *New Journal of Physics*, vol. 17, no. 2, p. 023037, 2015.
- [47] H. Choi, F. T. Chong, D. Englund, and Y. Ding, “Fault tolerant non-clifford state preparation for arbitrary rotations,” *arXiv preprint arXiv:2303.17380*, 2023.
- [48] C. Gidney, N. Shutty, and C. Jones, “Magic state cultivation: growing t states as cheap as cnot gates,” *arXiv preprint arXiv:2409.17595*, 2024.
- [49] C. Vuillot, L. Lao, B. Criger, C. G. Almudéver, K. Bertels, and B. M. Terhal, “Code deformation and lattice surgery are gauge fixing,” *New Journal of Physics*, vol. 21, no. 3, p. 033028, 2019.
- [50] K. Yin, X. Fang, T. S. Humble, A. Li, Y. Shi, and Y. Ding, “Surf-deformer: Mitigating dynamic defects on surface code via adaptive deformation,” 2024.
- [51] B. W. Reichardt, A. Paetznick, D. Aasen, I. Basov, J. M. Bello-Rivas, P. Bonderson, R. Chao, W. van Dam, M. B. Hastings, A. Paz *et al.*, “Logical computation demonstrated with a neutral atom quantum processor,” *arXiv preprint arXiv:2411.11822*, 2024.
- [52] L. Postler, S. Heußen, I. Pogorelov, M. Rispler, T. Feldker, M. Meth, C. D. Marciniak, R. Stricker, M. Ringbauer, R. Blatt *et al.*, “Demonstration of fault-tolerant universal quantum gate operations,” *Nature*, vol. 605, no. 7911, pp. 675–680, 2022.
- [53] B. W. Reichardt, D. Aasen, R. Chao, A. Chernoguzov, W. van Dam, J. P. Gaebler, D. Gresh, D. Lucchetti, M. Mills, S. A. Moses *et al.*, “Demonstration of quantum computation and error correction with a tesseract code,” *arXiv preprint arXiv:2409.04628*, 2024.
- [54] M. Kang, W. C. Campbell, and K. R. Brown, “Quantum error correction with metastable states of trapped ions using erasure conversion,” *PRX Quantum*, vol. 4, no. 2, p. 020358, 2023.
- [55] P. S. Rodriguez, J. M. Robinson, P. N. Jepsen, Z. He, C. Duckering, C. Zhao, K.-H. Wu, J. Campo, K. Bagnall, M. Kwon *et al.*, “Experimental demonstration of logical magic state distillation,” *arXiv preprint arXiv:2412.15165*, 2024.
- [56] Z. He, D. Amaro, R. Shaydulin, and M. Pistoia, “Performance of quantum approximate optimization with quantum error detection,” *arXiv preprint arXiv:2409.12104*, 2024.
- [57] R. Shaydulin, C. Li, S. Chakrabarti, M. DeCross, D. Herman, N. Kumar, J. Larson, D. Lykov, P. Minssen, Y. Sun *et al.*, “Evidence of scaling advantage for the quantum approximate optimization algorithm on a classically intractable problem,” *Science Advances*, vol. 10, no. 22, p. eadm6761, 2024.
- [58] S. Boulebnane and A. Montanaro, “Solving boolean satisfiability problems with the quantum approximate optimization algorithm,” *PRX Quantum*, vol. 5, no. 3, p. 030348, 2024.
- [59] S. Omanakuttan, Z. He, Z. Zhang, T. Hao, A. Babakhani, S. Boulebnane, S. Chakrabarti, D. Herman, J. Sullivan, M. A. Perlin *et al.*, “Threshold for fault-tolerant quantum advantage with the quantum approximate optimization algorithm,” *arXiv preprint arXiv:2504.01897*, 2025.
- [60] T. Hao, Z. He, R. Shaydulin, J. Larson, and M. Pistoia, “End-to-end protocol for high-quality qaoa parameters with few shots,” *arXiv preprint arXiv:2408.00557*, 2024.
- [61] Z. He, R. Shaydulin, S. Chakrabarti, D. Herman, C. Li, Y. Sun, and M. Pistoia, “Alignment between initial state and mixer improves qaoa performance for constrained optimization,” *npj Quantum Information*, vol. 9, no. 1, p. 121, 2023.
- [62] B. Bauer, S. Bravyi, M. Motta, and G. K.-L. Chan, “Quantum algorithms for quantum chemistry and quantum materials science,” *Chemical reviews*, vol. 120, no. 22, pp. 12 685–12 717, 2020.
- [63] N. J. Pearson, “Simulating many-body quantum systems: Quantum algorithms and experimental realisation,” Ph.D. dissertation, ETH Zurich, 2020.
- [64] M. E. Beverland, P. Murali, M. Troyer, K. M. Svore, T. Hoefler, V. Kliuchnikov, G. H. Low, M. Soeken, A. Sundaram, and A. Vashchillo, “Assessing requirements to scale to practical quantum advantage,” *arXiv preprint arXiv:2211.07629*, 2022.
- [65] K. More, “Vqe using qiskit,” <https://github.com/ketan-more-github/VQE-Using-Qiskit>, 2021, accessed: 2025-04-09.
- [66] S. P. Singh, “The ising model: Brief introduction and its application,” in *Solid state physics-metastable, spintronics materials and mechanics of deformable bodies-recent progress*. IntechOpen, 2020.
- [67] F. Bonechi, E. Celeghini, R. Giachetti, E. Sorace, and M. Tarlini, “Heisenberg xxz model and quantum galilei group,” *Journal of Physics A: Mathematical and General*, vol. 25, no. 15, p. L939, 1992.
- [68] D. Horsman, A. G. Fowler, S. Devitt, and R. Van Meter, “Surface code quantum computing by lattice surgery,” *New Journal of Physics*, vol. 14, no. 12, p. 123011, 2012.
- [69] O. Higgott, “Pymatching: A python package for decoding quantum codes with minimum-weight perfect matching,” *ACM Transactions on Quantum Computing*, vol. 3, no. 3, pp. 1–16, 2022.
- [70] S. B. Bravyi and A. Y. Kitaev, “Quantum codes on a lattice with boundary,” *arXiv preprint quant-ph/9811052*, 1998.
- [71] A. G. Fowler and C. Gidney, “Low overhead quantum computation using lattice surgery,” *arXiv preprint arXiv:1808.06709*, 2018.
- [72] N. J. Ross and P. Selinger, “Optimal ancilla-free clifford+ t approximation of z-rotations,” *arXiv preprint arXiv:1403.2975*, 2014.
- [73] E. Knill, “Fault-tolerant postselected quantum computation: Threshold analysis,” *arXiv preprint quant-ph/0404104*, 2004.
- [74] J. I. Cirac and P. Zoller, “Quantum computations with cold trapped ions,” *Physical review letters*, vol. 74, no. 20, p. 4091, 1995.
- [75] W. Paul, “Electromagnetic traps for charged and neutral particles,” *Reviews of modern physics*, vol. 62, no. 3, p. 531, 1990.

- [76] C. D. Bruzewicz, R. McConnell, J. Chiaverini, and J. M. Sage, “Scalable loading of a two-dimensional trapped-ion array,” *Nature communications*, vol. 7, no. 1, p. 13005, 2016.
- [77] K. Mølmer and A. Sørensen, “Multiparticle entanglement of hot trapped ions,” *Physical Review Letters*, vol. 82, no. 9, p. 1835, 1999.
- [78] W. Zhang, G. Tang, K. Liu, X. Yuan, Y. Shen, Y. Wu, and X.-M. Zhang, “Robust mølmer-sørensen gate against symmetric and asymmetric errors,” *Quantum Science and Technology*, 2025.
- [79] C. Ballance, T. Harty, N. Linke, and D. Lucas, “High-fidelity two-qubit quantum logic gates using trapped calcium-43 ions,” *arXiv preprint arXiv:1406.5473*, 2014.
- [80] R. Srinivas, S. Burd, H. Knaack, R. Sutherland, A. Kwiatkowski, S. Glancy, E. Knill, D. Wineland, D. Leibfried, A. C. Wilson *et al.*, *Nature*, vol. 597, no. 7875, pp. 209–213, 2021.
- [81] C. R. Clark, H. N. Tinkey, B. C. Sawyer, A. M. Meier, K. A. Burkhardt, C. M. Seck, C. M. Shappert, N. D. Guise, C. E. Volin, S. D. Fallek *et al.*, “High-fidelity bell-state preparation with ca+ 40 optical qubits,” *Physical Review Letters*, vol. 127, no. 13, p. 130505, 2021.
- [82] A. Myerson, D. Szwer, S. Webster, D. Allcock, M. Curtis, G. Imreh, J. Sherman, D. Stacey, A. Steane, and D. Lucas, “High-fidelity readout of trapped-ion qubits,” *Physical Review Letters*, vol. 100, no. 20, p. 200502, 2008.
- [83] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof, “Experimental issues in coherent quantum-state manipulation of trapped atomic ions,” *Journal of research of the National Institute of Standards and Technology*, vol. 103, no. 3, p. 259, 1998.
- [84] D. Kielpinski, C. Monroe, and D. J. Wineland, “Architecture for a large-scale ion-trap quantum computer,” *Nature*, vol. 417, no. 6890, pp. 709–711, 2002.
- [85] R. Blatt and D. Wineland, “Entangled states of trapped atomic ions,” *Nature*, vol. 453, no. 7198, pp. 1008–1015, 2008.
- [86] P. Murali, D. M. Debroy, K. R. Brown, and M. Martonosi, “Architecting noisy intermediate-scale trapped ion quantum computers,” in *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2020, pp. 529–542.
- [87] A. Ovide, D. Cuomo, and C. G. Almudever, “Scaling and assigning resources on ion trap qccd architectures,” in *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 1. IEEE, 2024, pp. 959–970.
- [88] C. Löschnauer, J. M. Toba, A. Hughes, S. King, M. Weber, R. Srinivas, R. Matt, R. Nourshargh, D. Allcock, C. Ballance *et al.*, “Scalable, high-fidelity all-electronic control of trapped-ion qubits,” *arXiv preprint arXiv:2407.07694*, 2024.
- [89] R. D. Delaney, L. R. Sletten, M. J. Cich, B. Estey, M. I. Fabrikant, D. Hayes, I. M. Hoffman, J. Hostetter, C. Langer, S. A. Moses *et al.*, “Scalable multispecies ion transport in a grid-based surface-electrode trap,” *Physical Review X*, vol. 14, no. 4, p. 041028, 2024.
- [90] M. Foss-Feig, G. Pagano, A. C. Potter, and N. Y. Yao, “Progress in trapped-ion quantum simulation,” *Annual Review of Condensed Matter Physics*, vol. 16, 2024.
- [91] “Quantinuum accelerates the path to Universal Fully Fault-Tolerant Quantum Computing,” 2024.
- [92] D. Poulin, “Stabilizer formalism for operator quantum error correction,” *Physical review letters*, vol. 95, no. 23, p. 230504, 2005.
- [93] O. Higgott and N. P. Breuckmann, “Subsystem codes with high thresholds by gauge fixing and reduced qubit overhead,” *Physical Review X*, vol. 11, no. 3, p. 031039, 2021.
- [94] A. deMarti iOlius, P. Fuentes, R. Orús, P. M. Crespo, and J. E. Martinez, “Decoding algorithms for surface codes,” *Quantum*, vol. 8, p. 1498, 2024.
- [95] T. LeBlond, R. S. Bennink, J. G. Lietz, and C. M. Seck, “Tisc: A surface code compiler and resource estimator for trapped-ion processors,” in *Proceedings of the SC’23 Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis*, 2023, pp. 1426–1435.
- [96] G. Li, Y. Ding, and Y. Xie, “Tackling the qubit mapping problem for nisq-era quantum devices,” in *Proceedings of the twenty-fourth international conference on architectural support for programming languages and operating systems*, 2019, pp. 1001–1014.
- [97] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. Allman, C. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer *et al.*, “Demonstration of the trapped-ion quantum ccd computer architecture,” *Nature*, vol. 592, no. 7853, pp. 209–213, 2021.
- [98] Quantinuum, “Quantinuum system model h1 product data sheet,” Jun. 2025, version 8.0.1, accessed July 31, 2025. [Online]. Available: https://docs.quantinuum.com/systems/data_sheets/Quantinuum%20H1%20Product%20Data%20Sheet.pdf
- [99] —, “Quantinuum system model h2 product data sheet,” Jun. 2025, version 1.0.0, accessed July 31, 2025. [Online]. Available: https://docs.quantinuum.com/systems/data_sheets/Quantinuum%20H2%20Product%20Data%20Sheet.pdf
- [100] S. Sarkar, M. Nambiar, S. G. Srinivasan *et al.*, “A modular quantum-classical framework for simulating chemical reaction pathways accurately,” *arXiv preprint arXiv:2210.08930*, 2022.
- [101] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, “The theory of variational hybrid quantum-classical algorithms,” *New Journal of Physics*, vol. 18, no. 2, p. 023023, 2016.
- [102] Qiskit Development Team, “Qiskit aer: High performance simulator for quantum circuits,” 2025, accessed: 2025-02-25. [Online]. Available: <https://github.com/Qiskit/qiskit-aer>
- [103] A. H. Wright, “Genetic algorithms for real parameter optimization,” in *Foundations of genetic algorithms*. Elsevier, 1991, vol. 1, pp. 205–218.
- [104] J. H. Holland, “Genetic algorithms,” *Scientific american*, vol. 267, no. 1, pp. 66–73, 1992.
- [105] C. Gidney, “Stim: a fast stabilizer circuit simulator,” *Quantum*, vol. 5, p. 497, 2021.