

SURGEQ: A Hybrid Framework for Ultra-Fast Quantum Processor Design and Crosstalk-Aware Circuit Execution

Xinxuan Chen^{*1}, Hongxiang Zhu^{*1}, Zhaohui Yang[†], Zhaofeng Su^{*‡}, Jianxin Chen[§], Feng Wu^{¶12}, Huihai Zhao^{¶12}

^{*}School of Computer Science and Technology, University of Science and Technology of China, Hefei 230026, China

[†]Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Hong Kong

[‡]Research Institute of Quantum Technology, The Hong Kong Polytechnic University, Hong Kong

[§]Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China

[¶]Zhongguancun Laboratory, Beijing 100194, China

Abstract—Executing quantum circuits on superconducting platforms requires balancing the trade-off between gate errors and crosstalk. To address this, we introduce SURGEQ, a hardware-software co-design strategy consisting of a design phase and an execution phase, to achieve accelerated circuit execution and improve overall program fidelity. SURGEQ employs coupling-strengthened, faster two-qubit gates while mitigating their increased crosstalk through a tailored scheduling strategy. With detailed consideration of composite noise models, we establish a systematic evaluation pipeline to identify the optimal coupling strength. Evaluations on a comprehensive suite of real-world benchmarks show that SURGEQ generally achieves higher fidelity than up-to-date baselines, and remains effective in combating exponential fidelity decay, achieving up to a million-fold improvement in large-scale circuits.

I. INTRODUCTION

To achieve high-fidelity execution of quantum applications, a quantum processor must adhere to specific design principles. In particular, errors related to the two-qubit gates have been identified as the primary concern in practical large-scale quantum computing [1], [2], [3]. Due to their sensitivity to noise-induced decoherence, quantum processors benefit from shorter two-qubit gate times through strengthened couplings between adjacent qubits. On the other hand, executing quantum circuits swiftly by using parallel operations is crucial to mitigate decoherence effects. However, fast operations and parallelism are sometimes mutually exclusive. Stronger qubit coupling exacerbates stray interactions, leading to severe crosstalk errors between simultaneous two-qubit gates and heightening the need for scheduling strategies to mitigate the spatio-temporal crosstalk. Balancing these two kinds of errors calls for a deliberate selection of coupling strength and gate time.

Crosstalk, as highlighted above, becomes particularly severe under stronger couplings and parallel two-qubit operations,

¹Co-first authors.

²Corresponding authors. zhaohuihai@iqubit.org, wufeng@iqubit.org.

Supported by Zhongguancun Laboratory, and sponsored by Quantum Science and Technology-National Science and Technology Major Project (Grant No. 2021ZD0302901)

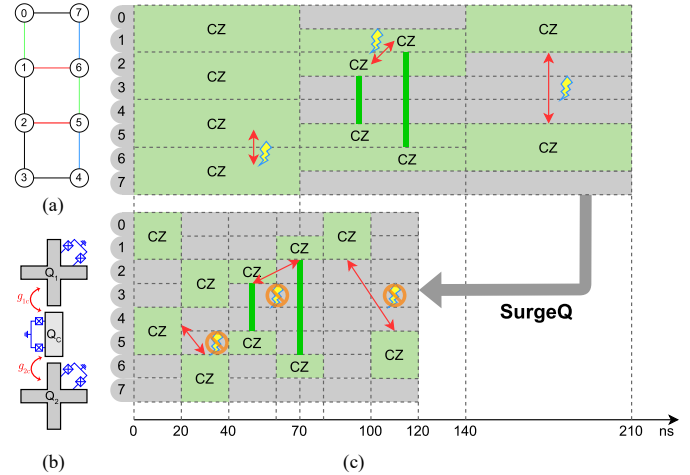


Fig. 1. Overview of SURGEQ's strategies on circuit execution. (a) Qubit connectivity graph for the example circuit. (b) Illustration of tunable couplers with adjustable coupling strengths. (c) Execution schedules before (top) and after (bottom) applying SURGEQ. The initial schedule contains multiple simultaneous CZ gates on adjacent qubits, leading to crosstalk events (highlighted with red arrows and yellow lightning markers, corresponding to colored coupling paths in (a)). The SURGEQ-optimized schedule jointly applies gate-level scheduling and coupling-strength tuning to resolve all crosstalk conflicts and significantly reduce execution time.

and remains one of the most significant challenges in executing large-scale circuits. To address this, several strategies have been introduced. For example, frequency allocation and adjustment techniques are employed to mitigate the detrimental effects of stray interactions caused by frequency crowding [4], [5]. Other solutions target different stages of the compilation process. [6] introduces a collaborative approach that integrates pulse optimization with scheduling, while [7] incorporates crosstalk-aware mapping strategies alongside gate scheduling. In contrast, [8] explores the trade-off between crosstalk mitigation and decoherence error reduction, employing optimization algorithms to minimize the impact of various errors within the

circuit.

Crosstalk is not unique to quantum computing; its classical counterparts have been extensively studied in areas like integrated circuits, wireless communication, and signal processing [9], [10], [11]. Inspired by the idea of crosstalk cancellation used to accelerate on-chip buses [9], we propose and develop SURGEQ—a hardware–software co-design framework implemented in two phases: a design stage that strengthens couplings to realize faster two-qubit gates, and an execution stage employing tailored scheduling to suppress the amplified crosstalk resulting from two-qubit gate acceleration. Unlike previous crosstalk-aware scheduling approaches relying on software-only mitigation [8], frequency allocation [4], or pulse-level co-optimization without design-stage considerations [6], our scheduler is explicitly tailored to co-design with hardware-level gate acceleration. SURGEQ first leverages stronger coupling to realize faster two-qubit gates, and then applies scheduling to suppress the amplified crosstalk, a capability not addressed in prior works [4], [6], [7], [8]. Fig. 1 illustrates the strategies employed by SURGEQ to improve circuit execution. It contrasts an initial gate execution scenario with the result of applying SURGEQ, which jointly applies gate-level scheduling and coupling-strength tuning to resolve crosstalk-inducing conflicts and reduce execution time. While the specific circuit used is simplified, the visual comparison conveys the key intuition behind SURGEQ. The process begins with increasing the qubit coupling strength to enable faster gates, which is critical for improving overall execution performance. However, stronger coupling significantly amplifies crosstalk, making naive speed-up detrimental to circuit fidelity. To address this, SURGEQ tightly integrates gate scheduling with coupling-strength calibration, jointly restructuring the circuit to eliminate crosstalk while preserving the benefits of faster gates. The result is a synergistic effect that surpasses what either technique can achieve in isolation.

Having established SURGEQ as our co-design framework, we now highlight the concrete contributions as follows:

1. We propose SURGEQ, a co-designed framework that significantly accelerates circuit execution by tuning coupling strength to maximize two-qubit gate speed, while effectively mitigating the exacerbated crosstalk effects through circuit optimization and scheduling.
2. We introduce a systematic evaluation pipeline that fully accounts for the dominant error sources to identify the feasible optimal coupling strength, and show that omitting any stage of our strategy markedly degrades overall performance.
3. We further validate SURGEQ’s effectiveness and versatility through a comprehensive evaluation on a diverse suite of real-world–inspired quantum circuits. Compared with other baselines, SURGEQ achieves higher fidelity in most cases and remains effective in combating exponential fidelity decay, reaching up to a million-fold improvements on larger-scale circuits.

II. BACKGROUND

Two-qubit gates present a greater challenge than single-qubit gates for quantum processors’ performance [12], [13], [1], [2], [3]. Issues such as stray interactions due to the weak anharmonicity of transmons [14], [15] and increased frequency crowding as systems scale up [16], [17] have been partially mitigated by tunable couplers. Such a design facilitates high-fidelity two-qubit gates in scalable processors.

Tunable couplers connect two qubits through electromagnetic interactions [18]. Specifically, the strength of this interaction can be adjusted by modifying the inductance and capacitance [19]. The Hamiltonian describing two transmon qubits connected by a transmon tunable coupler, as illustrated in Fig. 1(b), can be approximately given by (hereafter $\hbar = 1$) $H = \sum_{i=1,2,c} (\omega_i a_i^\dagger a_i + \frac{\alpha_i}{2} a_i^\dagger a_i^\dagger a_i a_i) + \sum_{i=1,2} g_{ic} (a_i + a_i^\dagger)(a_c + a_c^\dagger)$, where the subscript $i \in \{1, 2, c\}$ labels transmon Q_i . Here a_i^\dagger and a_i are the corresponding creation and annihilation operators, ω_i is the qubit frequency, α_i is the anharmonicity, and g_{ic} denotes the coupling strength between two qubits Q_i and Q_c . When the coupler is in its ground state, the ZZ interaction between Q_1 and Q_2 is defined as $\zeta_{ZZ} \equiv (E_{101} - E_{100}) - (E_{001} - E_{000})$, where E_{jkl} denotes the eigenenergy associated with the eigenstate $|jkl\rangle$ in the eigenbasis $|Q_1, C, Q_2\rangle$. The tunable coupler can be biased to tune up the coupling strength between $|101\rangle$ and $|002\rangle$, shifting E_{101} and activating the ZZ interaction, which scales as $g_{1c}g_{2c}$ [18], [20]. By accumulating this ZZ interaction over time, we can natively implement any gate from the CPhase(θ) family [14]. Usually the CZ = CPhase($\pi/2$) gate is calibrated on hardware to serve as the dominant two-qubit basis gate. Faster CZ operation can be realized via a larger ZZ interaction within the coupling strength bound.

However, there are non-negligible crosstalk errors among superconducting qubits arising from spatial and temporal correlations of quantum operations, especially for qubits with larger interactions [21]. In this work, we focus on the spatially non-local effects of gates, particularly the adjacent parallel two-qubit gates. In tunable coupler architectures, adjacent parallel two-qubit gates increase the likelihood of higher energy levels of couplers and qubits approaching each other, leading to significantly larger errors, compared to isolated two-qubit gates [5]. Such interaction creates unwanted ZZ terms across multiple qubits and couplers.

III. PROPOSED FRAMEWORK

A. Overview

Figure 2 depicts the workflow of the proposed SURGEQ framework. To fairly evaluate the trade-offs of accelerating quantum operations, we compare all scenarios using highly optimized quantum circuits. While circuit optimization is an active area of research, we emphasize simplicity and leverage third-party SDKs such as QISKIT to highlight our core contributions. A quantum circuit intended for execution on a chip is processed through a compilation workflow borrowed from third-party SDKs, referred to as *base compilation*. This

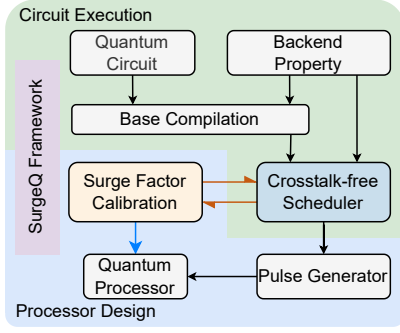


Fig. 2. Overview of the SURGEQ framework. In addition to the standard workflow of quantum processor design and circuit execution, the Surge Factor Calibration Unit is introduced to evaluate performance gains achieved through highly optimized circuit execution that avoids crosstalk penalties, as indicated by the bi-directional red arrows. The optimal Surge factor guides the selection of appropriate coupling strengths, as shown by the blue arrow.

workflow incorporates the backend properties of the chip, such as qubit connectivity topology, and outputs a compatible optimized circuit. We then introduce a *Crosstalk-Free Scheduler* that prevents parallel operations on adjacent qubit pairs, thereby eliminating nearest-neighbor ZZ crosstalk at the cost of increasing circuit depth. Unlike [8], our scheduler does not rely on physical characterization data and can reduce the depth overhead of the serialization process.

We introduce a *Surge factor*, s , which indicates the amplification of the coupling strength during the processor design phase to balance errors from different sources. Specifically, g_{1c} and g_{2c} are scaled to s times their respective reference values. It is optimized in the *Surge Factor Calibration* module to determine the optimal s that maximizes circuit fidelity.

In the following subsections, we will provide a detailed discussion of both modules.

B. Surge Factor Calibration

The Surge Factor Calibration module begins by extracting hardware parameters, such as qubit connectivity, gate durations, decoherence time and crosstalk, from specifications or standard characterization experiments on a testing quantum processor. Using these parameters, this module estimates errors of this processor, which is provided to a scheduler that performs crosstalk-aware gate scheduling, simulates circuit execution to compute the fidelity across a range of surge factors. Then this module identifies an optimal surge factor to be applied in a new processor design intended to be fabricated as the final hardware to execute the circuit.

The physical implementation of the design depends on the processor architecture. For fixed coupling architectures, tuning the coupling between qubits is achieved by adjusting the geometry (i.e., size and shape) of qubits. In this work, we take tunable coupling architecture as an example, of which the geometry of both qubits and couplers requires modification.

To evaluate the effect of the Surge Factor s on quantum gate performance, this module employs a comprehensive error

Algorithm 1: Crosstalk-Free Scheduling

Input: A quantum circuit with timing information

Output: A scheduled DAG without ZZ crosstalk

```

1  $\mathcal{D} = \{V_D, E_D\} \leftarrow$  translate input circuit into DAG
2  $\mathcal{T} = \{V_T, E_T\} \leftarrow$  initialize an empty DAG
3 while  $\mathcal{D} \neq \emptyset$  do
4    $V_0 \leftarrow$  head nodes of  $\mathcal{D}$ 
5    $S(C) \leftarrow$  nodes of single-(two-)qubit gates in  $V_0$ 
6    $M \leftarrow$  MISGATES( $C$ )
7    $R = C - M$ 
8    $R' \leftarrow$  PATTERNSUBSTITUTE( $R, \mathcal{D}$ )
9   Layer  $\leftarrow S \cup M \cup R'$ 
10   $\mathcal{T}$ .add(Layer)
11   $\mathcal{D}$ .delete(Layer)
12 return  $\mathcal{T}$ 
13 Procedure MISGATES(Node set  $C$ )
14    $\mathcal{G} \leftarrow$  init an empty undirected graph  $\mathcal{G}(V, E)$ 
15    $\mathcal{G}.V \leftarrow C$ 
16   foreach  $\{v_1, v_2\} \subset \mathcal{G}.V$  do
17     if  $v_1$  and  $v_2$  have nearest ZZ crosstalk
18       then
19          $\mathcal{G}.E.add((v_1, v_2))$ 
20    $M \leftarrow$  nodes that constitute the MIS of  $\mathcal{G}$ 
21   return  $M$ 
22 Procedure PATTERNSUBSTITUTE(Nodes  $V$ , DAG  $\mathcal{D}$ )
23   foreach  $v \in V$  do
24     if  $v$  and its successors in  $\mathcal{D}$  match a pattern
25       then
26         Substitute the pattern in  $\mathcal{D}$ 
27          $s_1, s_2 \leftarrow$  the new head nodes
28          $R'.add(s_1, s_2)$ 
29   return  $R'$ 

```

model. In this model, the relationships between different error sources and the Surge Factor s are listed below.

- 1) **Two-Qubit Decoherence Error:** The two-qubit gate error can be attributed to multiple sources [22]. First, we assume the two-qubit gate time t_g is inversely proportional to the qubit-qubit ZZ interaction strength. Based on the discussion in the previous section, this implies $s^2 \propto 1/t_g$. We decompose the decoherence error into two components: E_1 and E_ϕ . The relaxation error E_1 , arising from qubit relaxation (T_1), follows the relation $E_1 = \frac{1}{3}t_g/T_1$, which is proportional to the gate time, yielding $E_1 \propto t_g \propto s^{-2}$. For pure dephasing, assumed to be dominated by $1/f$ flux noise, the error is given by $E_\phi = \frac{1}{3}t_g^2/T_\phi^2$, proportional to the square of the gate time, leading to $E_\phi \propto t_g^2 \propto s^{-4}$.
- 2) **Two-Qubit Coherent Error:** The dominant coherent error in the two-qubit gate, the non-adiabatic error E_A , is determined by the adiabatic condition factor and the pulse shape [22]. However, there is no simple relationship between s and E_A . Considering that varying s affects E_A through two opposing mechanisms—larger energy level splitting from stronger coupling and faster energy changes from shorter gate times—which can partially offset each other, we assume E_A remains constant

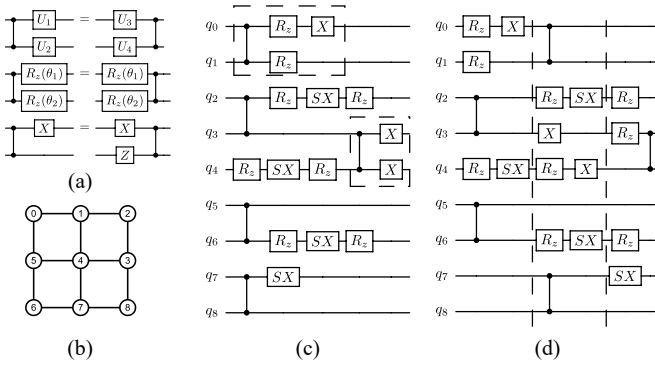


Fig. 3. (a) Pattern substitution demonstration. The topmost pattern represents the desired form of pattern substitution. The rest patterns illustrate the current set of pattern substitutions available in our library. (b) Example topological layout of the qubits for (c)(d). The nodes in the topological layout represent qubits while edges denote the coupling between qubits. (c)(d) An example of a scheduling algorithm for the quantum circuit (c) with result (d). We omit the rotation angle θ around the Z axis of $R_z(\theta)$ gate, as the pattern substitution and merging of $R_z(\theta)$ gates involving changes in θ of certain $R_z(\theta)$ are irrelevant to our approach. The duration of the whole layer is decided by the longest gate inside, with the timing inconsistency compensated by idle time in each layer. The enclosed parts in (c) are identified as structures that can undergo pattern substitution. In the first loop of the algorithm, the two-qubit gate set $\{q_2-q_3, q_5-q_6\}$ is the MIS found, while a pattern containing q_0-q_1 , along with their subsequent single-qubit gates, has been identified. The same line of thought is also applied in the following loop.

for different s in this work.

- 3) **Single-Qubit Error:** On the other hand, the coupling strength does not affect single-qubit gates. Thus, the decoherence errors E_1 , E_ϕ , and the coherent error E_c of single-qubit gates remain constant irrespective of s .
- 4) **Crosstalk Effects on Single-Qubit Errors:** Crosstalk errors arise when adjacent qubits execute single-qubit or two-qubit gates simultaneously [5]. For single-qubit gates, crosstalk manifests as cross-drive on nearby qubits, leading to an error E_{1Q} . This error scales with the square of the cross-drive strength, which is proportional to the qubit-qubit coupling strength, resulting in $E_{1Q} \propto s^4$.
- 5) **Crosstalk Effects on Two-Qubit Errors:** For parallel two-qubit gates, the error E_{2Q} originates from stray ZZ interactions involving higher energy levels of four qubits, such as $|0002\rangle$ and $|1000\rangle$. These interactions scale as s^6 , while the gate time scales as s^{-2} , leading to $E_{2Q} \propto s^4$.

C. Crosstalk-Free Scheduler

Our scheduler staggers adjacent two-qubit gates while minimizing additional circuit depth. We model circuits as Directed Acyclic Graphs (DAGs), where nodes represent gates and edges represent dependencies. We distinguish the layer structure by inserting barrier nodes between consecutive layers, thereby dividing the DAG into layers and enabling sequential scheduling.

Algorithm 1 converts the input circuit into a DAG D and iteratively extracts head nodes V_0 . Among two-qubit gates

$C \subseteq V_0$, we apply a Maximum Independent Set (MIS) algorithm to select the largest subset M free of nearest-neighbor ZZ crosstalk. Remaining gates R are checked for predefined substitution patterns, producing modified nodes R' . Together with single-qubit gates S , these (M, R') form a scheduled layer, appended to the target DAG T . The process repeats until D is empty, with all $R_z(\theta)$ rotations excluded from the layer count (e.g., “ $-R_z(\theta) - SX-$ ” is treated as one layer). The final DAG fully eliminates nearest-neighbor ZZ crosstalk. An example is illustrated in Fig. 3(b)(c)(d).

Within the MIS algorithm, two-qubit gates are mapped to an undirected graph where edges indicate potential crosstalk. The MIS gives the maximum parallel set, though deferred gates increase idle time and circuit depth. To mitigate this, we apply pattern substitution: if some deferred gates and their successors match a predefined template, a substitution is applied. This substitution reduces the overhead introduced by MIS, and as the pattern library expands, the scheduler can further minimize depth growth. The desired form of pattern substitution is illustrated at the top of Fig. 3(a), where the rest illustrates the current set of pattern substitutions available in our library.

IV. EVALUATION

A. Experimental Setup

1) **Digital error model:** Digital quantum computing relies on the assumption that errors are discrete, localized, and independent both temporally and spatially. This assumption was experimentally validated in [16], where a digital error model was also proposed to characterize the expected behavior of large-scale quantum circuits.

Error models for individual gates are thoroughly discussed in Subsec. III-B. The digital error model describes the expected behavior of a circuit composed of these individual gates, focusing on accuracy or fidelity within the context of quantum computing. Under the assumption of the digital error model, the circuit fidelity is calculated as the product of the fidelities of individual gates.

We note that during the execution of a two-qubit gate, one qubit is tuned away from its idle point, resulting in a significant reduction in its effective T_1 and T_ϕ due to stronger interaction with the coupler and increased flux noise [22]. In the equations below, the superscript “eff” denotes the qubit in this condition, while “idle” is used to represent the qubit in its idle state.

The fidelity of a single-qubit gate is given by $F_{1Q} = 1 - E_c - E_1^{\text{idle}} - E_\phi^{\text{idle}} - N_1 E_{1Q}$, where N_1 is the number of single-qubit gates adjacent to this single-qubit gate. Similarly, the fidelity of a two-qubit gate is given by $F_{2Q} = 1 - E_1^{\text{eff}} - E_1^{\text{idle}} - E_\phi^{\text{eff}} - E_\phi^{\text{idle}} - E_A - N_1 E_{1Q} - N_2 E_{2Q}$. Here, N_1 and N_2 count the single- and two-qubit gates adjacent to this two-qubit gate. When the qubit is idle, the fidelity decreases as $F_{\text{idle}} = 1 - E_1^{\text{idle}} - E_\phi^{\text{idle}}$. Finally, the circuit fidelity is calculated by $F = \prod_{\{1Q \text{ gate}\}} F_{1Q} \prod_{\{2Q \text{ gate}\}} F_{2Q} \prod_{\{\text{idle time}\}} F_{\text{idle}}$.

2) **Setup:** Our compilation experiments are conducted on an Intel Core i5-14500 processor running Linux 6.8.0. We implement our approach in Python 3.9, utilizing QISKIT 1.0.2 [23]

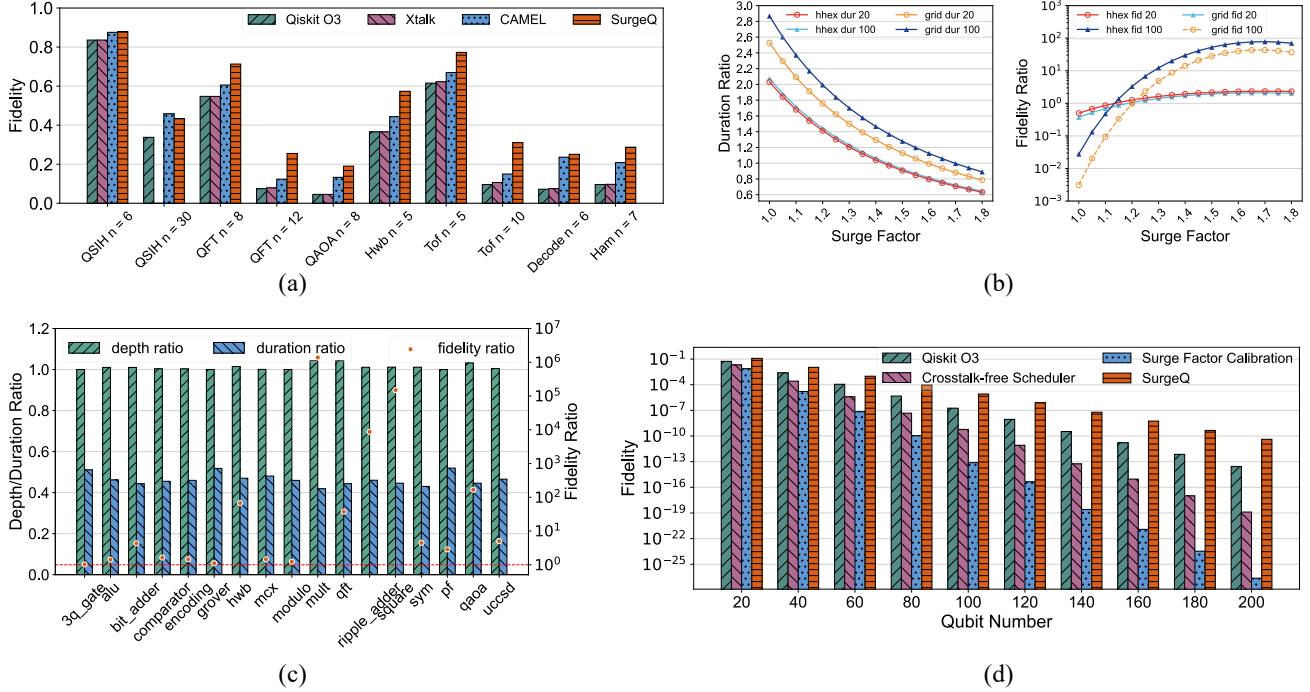


Fig. 4. (a) Comparison of fidelity across several representative circuit families under different baselines. All circuits are from the dataset described in IV-A4, and n denotes the number of qubits, with QSIH circuits fixed to 2 Trotter steps. Due to XtalkSched’s exponential-time complexity and the resulting prohibitive compilation times, results for QSIH with $n=30$ are unavailable. Overall, SurgeQ achieves higher fidelity than the alternatives in most cases. (b) The Surge Factor Calibration for QSIH circuits. All the data are the ratio compared with the results obtained by using only the base compilation—QISKIT O3. A lower duration ratio and a higher fidelity ratio are indicative of superior performance. Legends such as hhhex/grid, dur/fid, and 20/100 indicate the corresponding data points’ attributes: topology (heavy hexagon or grid), metric (duration ratio or fidelity ratio), and circuit size (20-qubit or 100-qubit). (c) Benchmarking results of SURGEQ across a comprehensive suite of quantum circuits derived from real-world quantum applications, evaluated in terms of depth, duration, and fidelity ratios, with the Surge factor set to $s = 1.8$. (d) Ablation experiment of SURGEQ for QSIH execution on grid backend with $s = 1.7$

TABLE I
PHYSICAL PARAMETERS REALIZABLE WITH CURRENT HARDWARE.

T_1^{eff}	T_ϕ^{eff}	T_1^{idle}	T_ϕ^{idle}
40000 ns	1100 ns	50000 ns	40000 ns
E_c	E_A	E_{1Q}	E_{2Q}
3.25×10^{-6}	1.2×10^{-3}	2.13×10^{-5}	1.75×10^{-4}

and NetworkX 3.2.1 [24]. The base compilation used in our study is QISKIT optimization level 3 (-O3), which represents the highest optimization level for the original circuit.

3) *Quantum Backend*: We compile quantum circuits into a basic gate set $\{I, X, SX, R_Z(\theta), CZ\}$, where $SX = \sqrt{X}$, $R_Z(\theta) = e^{-i\theta Z/2}$. Referring to [5], we set the time of single-qubit gates S and SX to 20 ns, and initialize the CZ gate with a duration of 70 ns without adjusting the coupling strength. According to the Virtual Z technology [25], we set the time of $R_Z(\theta)$ as 0 for all θ available. The backends we are using include a 127-qubit heavy-hex (hhhex in brief) structure and a 225-qubit grid structure.

4) *Benchmarks*: To conduct thorough benchmarking, we carefully select diverse categories of quantum programs, spanning 17 categories and a total of 127 benchmarks. Most programs are drawn from established benchmark suites [26], [27], [28], with the exception of the *ripple_adder* and *qaoa* cate-

gories, which are generated by us. All the categories include *3q_gate*, *alu* (Arithmetic Logic Unit), *bit_adder*, *comparator*, *encoding* (encoding function), *grover*, *hwb* (Hidden Weighted Bit function), *mcx* (Multi-controlled-X gates), *modulo*, *mult* (Galois Field Multiplier), *qft* (Quantum Fourier Transform), *ripple_adder* (Ripple Carry Adder), *square* (Square Root function), *sym* (Symmetric function), *pf* (1D Heisenberg Product Formula), *qaoa*, and *uccsd* (UCCSD programs generated using QISKIT and PYSCF [29] from TKET’s published benchmark suite). This diverse selection provides a comprehensive evaluation of SURGEQ’s performance across a broad spectrum of quantum applications. We also select the second-order Trotterization of the Isotropic Heisenberg Hamiltonian evolution with open boundary conditions (referred to as QSIH) from [30], due to its numerous adjacent parallel operations that may result in significant overhead for the Crosstalk-Free Scheduler.

5) *Physical Parameters*: Building on studies of two-qubit gates and crosstalk in devices utilizing tunable couplers [22], [5], [31], the parameters in Table I reflect physical parameters achievable with current hardware and are used to calculate the fidelities of single-qubit gates, two-qubit gates, and idling.

B. Experimental Results

We will now systematically evaluate SURGEQ using the experimental setup described above. Since our scheduling

algorithm (Alg. 1) is heuristic and relies on a NetworkX approximation algorithm, we take the result with the minimum depth observed over 20 independent runs.

1) *Baselines*: We evaluate the efficacy of our framework by comparing it with several up-to-date baselines. We use QISKIT O3 as a crosstalk-agnostic optimizer and XtalkSched [8] and CAMEL [7] as crosstalk-aware baselines. For compatibility with CAMEL’s scheduling scheme, we set the CZ gate time for it alone to 40 ns. As shown in Fig. 4(a), across diverse circuit families, SurgeQ generally attains higher fidelity than all baselines. A notable exception occurs on 30-qubit QSIH circuits, where crosstalk is particularly severe: CAMEL’s compensation pulse scheme affords a greater reduction in circuit depth, yielding fidelity slightly higher than that of SurgeQ.

2) *Feasibility*: SURGEQ offers both advantages and disadvantages: while accelerating operations to reduce decoherence errors, the increased circuit depth caused by the Crosstalk-Free Scheduler may partially offset these benefits, thus lessening the overall performance impact. Therefore, we evaluate its feasibility using QSIH, the highly dense circuit family mentioned before, as a worst-case test scenario.

We vary the Surge factor s within $[1.0, 1.8]$ to accelerate gate operations during crosstalk-aware scheduling of QSIH circuits. Values above 1.8 are excluded as they are not reliably attainable on current hardware and can introduce confounding effects such as pulse distortion [5], [20], [22], [16]. The performance gain is defined in terms of the duration ratio and fidelity ratio, calculated as the duration and fidelity of the resulting circuits of SURGEQ divided by those of QISKIT O3. The evaluation is conducted on the QSIH circuit family with 10 Trotter steps, considering different qubit counts (20 and 100) across two quantum chip topologies: heavy hexagonal (hhex) and grid.

Fig. 4(b) illustrates the Surge Factor Calibration, showing $s = 1.7$ as optimal for QSIH circuits. As s increases, the fidelity ratio improves in the QSIH algorithm with 20 qubits and grows even more rapidly with 100 qubits when $1 \leq s \leq 1.7$. However, a slight decrease occurs when $s \geq 1.7$, indicating that at this point, the increase in s significantly elevates E_{1Q} , offsetting error reductions from other sources. Notably, fidelity ratios increase faster in larger-scale circuits, outperforming QISKIT O3 by several orders of magnitude. For instance, when $s = 1.7$ and the qubit count is 100, fidelity ratios approach 100. The depth overhead of the crosstalk-free-scheduled circuit is $1.61\times$ that of the non-crosstalk-free-scheduled version. In all experiments, with qubit counts ranging from 20 to 200, this overhead consistently stays below $1.7\times$.

3) *Applicability*: We will further evaluate the applicability of SURGEQ. Given that the QSIH is chosen as the worst-case test scenario, where we still observe a significant performance boost, we anticipate that SURGEQ will demonstrate advantages across a wide range of real-world applications.

Our theoretical analysis in Subsec. III-B reveals that only E_{1Q} and E_{2Q} increase exponentially with s , while our Crosstalk-Free Scheduler eliminates E_{2Q} . Consequently, crosstalk errors from single-qubit gates (E_{1Q}) have become a

key factor limiting further efficiency improvements. However, the density of E_{1Q} errors varies significantly across different circuits. It is therefore not surprising that the Surge Factor Calibration selects $s = 1.7$ for the highly dense QSIH, while for most of the less dense real-world applications, it selects $s = 1.8$.

Figure 4(c) illustrates the benchmarking results for a comprehensive suite of representative circuits from real-world applications, where the result for each circuit type is the geometric mean of the results from circuits of the same type at different scales. These circuits utilize fewer parallel two-qubit gates on adjacent qubits, leading to lower depth and duration overhead for crosstalk-free scheduling. The duration ratio is a more appropriate metric in this context. For representativeness, some benchmarks include extremely small-scale circuits, such as 3-qubit circuits. All benchmarks show consistent duration improvements. Small-scale circuits still achieve significant fidelity improvements, though less apparent on a logarithmic scale. With increased system size, fidelity improvements can reach as high as 10^6 . These results demonstrate SURGEQ’s exceptional performance across diverse real-world applications.

4) *Ablation Study*: We further conduct an experiment to evaluate the impact of the two key modules in SURGEQ—the Surge Factor Calibration and the Crosstalk-Free Scheduler. For the Surge Factor Calibration, we amplify the coupling strength to accelerate quantum operations but execute the circuits without applying the Crosstalk-Free Scheduler, resulting in crosstalk penalties. For the Crosstalk-Free Scheduler, we keep the chip unchanged and focus solely on mitigating crosstalk effects through scheduling. Fig. 4(d) clearly demonstrates that the absence of any individual component in our strategy significantly degrades the overall performance.

5) *Scalability*: The time complexity of the crosstalk-free scheduling algorithm is $O(\frac{m^2}{n})$, with m and n denoting the gate count and the qubit count, respectively. Therefore, the core functionalities of SURGEQ will not become a bottleneck for scaling up. The SURGEQ framework leverages third-party compiler workflows, such as QISKIT at optimization level O3. Gate-level optimizations have a complexity ranging from $O(m)$ to $O(m^2)$, while qubit mapping is $O(n^2)$. Thus, SURGEQ is well-suited for scaling to systems with thousands of qubits.

V. SUMMARY AND OUTLOOK

SURGEQ presents a robust solution to the critical challenge of maintaining fidelity in large-scale quantum circuits. By synergizing coupling-strengthened processor design with ultra-fast two-qubit gates and crosstalk-aware circuit scheduling, our co-design framework effectively mitigates the exponential decay of program fidelity. By providing an optimized and practical design for both speed and noise, SURGEQ paves the way for realizing more complex and impactful quantum subroutines, from fault-tolerant quantum memories to application-specific modules like the Quantum Fourier Transform, advancing the pursuit of practical quantum advantage.

REFERENCES

- [1] Z. Chen, K. J. Satzinger, J. Atalaya, A. N. Korotkov, A. Dunsworth, D. Sank *et al.*, “Exponential suppression of bit or phase errors with cyclic error correction,” *Nature*, vol. 595, pp. 383–387, Jul. 2021.
- [2] X. Mi, P. Roushan, C. Quintana, S. Mandra, J. Marshall, C. Neill *et al.*, “Information Scrambling in Computationally Complex Quantum Circuits,” *Science*, vol. 374, pp. 1479–1483, Dec. 2021.
- [3] R. Acharya, I. Aleiner, R. Allen, T. I. Andersen, M. Ansmann, F. Arute *et al.*, “Suppressing quantum errors by scaling a surface code logical qubit,” *Nature*, vol. 614, pp. 676–681, Feb. 2023.
- [4] Y. Ding, P. Gokhale, S. F. Lin, R. Rines, T. Propson, and F. T. Chong, “Systematic crosstalk mitigation for superconducting qubits via frequency-aware compilation,” in *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture*. IEEE, 2020, pp. 201–214.
- [5] P. Zhao, K. Linghu, Z. Li, P. Xu, R. Wang, G. Xue *et al.*, “Quantum crosstalk analysis for simultaneous gate operations on superconducting qubits,” *PRX Quantum*, vol. 3, p. 020301, 2022.
- [6] L. Xie, J. Zhai, Z. Zhang, J. Allcock, S. Zhang, and Y.-C. Zheng, “Suppressing zz crosstalk of quantum computers through pulse and scheduling co-optimization,” in *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, 2022, pp. 499–513.
- [7] B.-H. Lu, P. Wang, Z.-Y. Chen, H.-Y. Liu, T.-P. Sun, P. Duan *et al.*, “Camel: Physically inspired crosstalk-aware mapping and gate scheduling for frequency-tunable quantum chips,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 44, no. 5, pp. 1968–1980, 2025.
- [8] P. Murali, D. C. McKay, M. Martonosi, and A. Javadi-Abhari, “Software mitigation of crosstalk on noisy intermediate-scale quantum computers,” in *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2020, pp. 1001–1016.
- [9] C. Duan and S. Khatri, “Exploiting crosstalk to speed up on-chip buses,” in *Proceedings Design, Automation and Test in Europe Conference and Exhibition*, vol. 2, 2004, pp. 778–783.
- [10] S. Sinha, R. Kar, and A. Bhattacharjee, “Bus encoding technique using forbidden transition free algorithm for cross-talk reduction for on-chip vlsi interconnect,” in *2010 International Conference on Advances in Computer Engineering*. IEEE, Jun. 2010, pp. 256–258.
- [11] M. Daraban and D. Pitica, “Forbidden switching patterns to mitigate crosstalk on parallel data buses — codec design,” in *2012 IEEE 18th International Symposium for Design and Technology in Electronic Packaging*, 2012, pp. 31–36.
- [12] D. M. Zajac, A. J. Sigillito, M. Russ, F. Borjans, J. M. Taylor, G. Burkard *et al.*, “Resonantly driven cnot gate for electron spins,” *Science*, vol. 359, pp. 439–442, 2018.
- [13] K. Wright, K. M. Beck, S. Debnath, J. Amini, Y. Nam, N. Grzesiak *et al.*, “Benchmarking an 11-qubit quantum computer,” *Nature Communications*, vol. 10, p. 5464, 2019.
- [14] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver, “A quantum engineer’s guide to superconducting qubits,” *Applied Physics Reviews*, vol. 6, p. 021318, 2019.
- [15] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, “Circuit quantum electrodynamics,” *Reviews of Modern Physics*, vol. 93, p. 025005, 2021.
- [16] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends *et al.*, “Quantum supremacy using a programmable superconducting processor,” *Nature*, vol. 574, pp. 505–510, 2019.
- [17] X. Han, T. Cai, X. Li, Y. Wu, Y. Ma, Y. Ma *et al.*, “Error analysis in suppression of unwanted qubit interactions for a parametric gate in a tunable superconducting circuit,” *Physical Review A*, vol. 102, p. 022619, 2020.
- [18] F. Yan, P. Krantz, Y. Sung, M. Kjaergaard, D. L. Campbell, T. P. Orlando *et al.*, “Tunable coupling scheme for implementing high-fidelity two-qubit gates,” *Physical Review Applied*, vol. 10, p. 054062, 2018.
- [19] Y. Chen, C. Neill, P. Roushan, N. Leung, M. Fang, R. Barends *et al.*, “Qubit architecture with high coherence and fast tunable coupling,” *Physical Review Letters*, vol. 113, p. 220502, 2014.
- [20] Y. Sung, L. Ding, J. Braumüller, A. Vepsäläinen, B. Kannan, M. Kjaergaard *et al.*, “Realization of high-fidelity cz and zz-free iswap gates with a tunable coupler,” *Physical Review X*, vol. 11, no. 2, p. 021058, 2021.
- [21] M. Sarovar, T. Proctor, K. Rudinger, K. Young, E. Nielsen, and R. Blume-Kohout, “Detecting crosstalk errors in quantum information processors,” *Quantum*, vol. 4, p. 321, Sep. 2020.
- [22] Y. Xu, J. Chu, J. Yuan, J. Qiu, Y. Zhou, L. Zhang *et al.*, “High-fidelity, high-scalability two-qubit gate scheme for superconducting qubits,” *Physical Review Letters*, vol. 125, p. 240503, 2020.
- [23] A. Javadi-Abhari, M. Treinish, K. Krsulich, C. J. Wood, J. Lishman, J. Gacon *et al.*, “Quantum computing with Qiskit,” 2024, arXiv:2405.08810.
- [24] A. A. Hagberg, D. A. Schult, and P. J. Swart, “Exploring network structure, dynamics, and function using networkx,” in *Proceedings of the 7th Python in Science Conference*, 2008, pp. 11–15.
- [25] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, “Efficient z gates for quantum computing,” *Physical Review A*, vol. 96, p. 022330, 2017.
- [26] R. Wille, D. Große, L. Teuber, G. W. Dueck, and R. Drechsler, “Revlib: An online resource for reversible functions and reversible circuits,” in *38th International Symposium on Multiple Valued Logic*. IEEE, 2008, pp. 220–225.
- [27] S. Sivarajah, S. Dilkes, A. Cowtan, W. Simmons, A. Edgington, and R. Duncan, “ $t|ket\rangle$: a retargetable compiler for nisq devices,” *Quantum Science and Technology*, vol. 6, p. 014003, 2020.
- [28] A. Zulehner, A. Paler, and R. Wille, “An efficient methodology for mapping quantum circuits to the ibm qx architectures,” *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 38, pp. 1226–1236, 2018.
- [29] Q. Sun, T. C. Berkelbach, N. S. Blunt, G. H. Booth, S. Guo, Z. Li *et al.*, “Pyscf: the python-based simulations of chemistry framework,” *Wiley Interdisciplinary Reviews: Computational Molecular Science*, vol. 8, p. e1340, 2018.
- [30] T. A. Chowdhury, K. Yu, M. A. Shamim, M. Kabir, and R. S. Sufian, “Enhancing quantum utility: simulating large-scale quantum spin chains on superconducting quantum computers,” *Physical Review Research*, vol. 6, p. 033107, 2024.
- [31] C. Wang, X. Li, H. Xu, Z. Li, J. Wang, Z. Yang *et al.*, “Towards practical quantum computers: Transmon qubit with a lifetime approaching 0.5 milliseconds,” *npj Quantum Information*, vol. 8, p. 3, 2022.