

Dependency-Aware Compilation for Surface Code Quantum Architectures

ABTIN MOLAVI, University of Wisconsin-Madison, USA

AMANDA XU, University of Wisconsin-Madison, USA

SWAMIT TANNU, University of Wisconsin-Madison, USA

AWS ALBARGHOUTHI, University of Wisconsin-Madison, USA

Practical applications of quantum computing depend on fault-tolerant devices with error correction. Today, the most promising approach is a class of error-correcting codes called *surface codes*. We study the problem of compiling quantum circuits for quantum computers implementing surface codes. Optimal or near-optimal compilation is critical for both efficiency and correctness. The compilation problem requires (1) *mapping* circuit qubits to the device qubits and (2) *routing* execution paths between interacting qubits. We solve this problem efficiently and near-optimally with a novel algorithm that exploits the *dependency structure* of circuit operations to formulate discrete optimization problems that can be approximated via *simulated annealing*, a classic and simple algorithm. Our extensive evaluation shows that our approach is powerful and flexible for compiling realistic workloads.

CCS Concepts: • **Software and its engineering** → **Compilers**; • **Hardware** → **Quantum computation**.

Additional Key Words and Phrases: quantum error-correction, simulated annealing

ACM Reference Format:

Abtin Molavi, Amanda Xu, Swamit Tannu, and Aws Albarghouthi. 2025. Dependency-Aware Compilation for Surface Code Quantum Architectures. *Proc. ACM Program. Lang.* 9, OOPSLA1, Article 82 (April 2025), 28 pages. <https://doi.org/10.1145/3720416>

1 Introduction

Quantum computation promises to surpass classical methods in important domains, potentially unlocking breakthroughs in materials science, chemistry, machine learning, and beyond. However, as individual physical qubits and operations are error-prone, these applications require an error-correction scheme for detecting and correcting faults. Quantum error-correction suppresses errors with redundancy: encoding the state of a single logical qubit using several physical qubits. Experimentalists have recently demonstrated error suppression for a single logical qubit [26, 52, 66] and small multi-qubit systems [11, 20, 27, 51].

To harness the full power of the fault-tolerant quantum computers on the horizon, we need optimizing compilers that convert circuit-level descriptions of quantum programs to error-corrected elementary operations while preserving as much parallelism as possible. Quantum compute is a scarce resource, so inefficient compilation can be extremely costly. Further, the longer the computation, the higher the probability of logical errors, which affect the result.

Therefore, our goal is to answer the following question:

How can we compile a given circuit for a fault-tolerant device such that execution time is minimized?

Authors' Contact Information: [Abtin Molavi](#), University of Wisconsin-Madison, USA, amolavi@wisc.edu; [Amanda Xu](#), University of Wisconsin-Madison, USA, axu44@wisc.edu; [Swamit Tannu](#), University of Wisconsin-Madison, USA, stannu@wisc.edu; [Aws Albarghouthi](#), University of Wisconsin-Madison, USA, aws@cs.wisc.edu.



This work is licensed under a [Creative Commons Attribution 4.0 International License](#).

© 2025 Copyright held by the owner/author(s).

ACM 2475-1421/2025/4-ART82

<https://doi.org/10.1145/3720416>

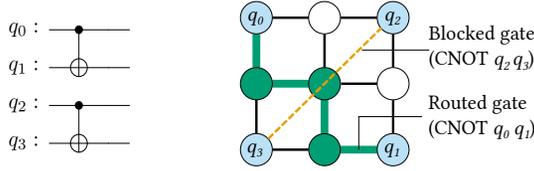


Fig. 1. Spatial constraints preventing parallel execution

Table 1. Feature summary of our approach and prior work

Compiler	Optimizing?	T gates?	Any arch?
LSC [62]	✗	✓	✓
EDPC [8]	✓* (route only)	✓	✗
Autobraid [31]	✓	✗	✗
DASCOT (this paper)	✓	✓	✓

We target a well-studied type of error-correction scheme called a *surface code* [24, 36, 43]. A surface code quantum device embeds logical qubits into a two-dimensional grid of physical qubits. Two-qubit gates impose limitations on the execution of a quantum circuit by introducing contention constraints. Each two-qubit gate occupies a path on the grid and simultaneous paths cannot cross. Gates which can theoretically be executed in parallel may be forced into sequential execution if the path of one “blocks” the other, as shown in Fig. 1. A compiler must carefully *map* qubits to grid locations and *route* two-qubit gates such that conflicts between gates are minimized and parallelism is maximized. We call this the *surface code mapping and routing* (SCMR) problem.

Existing work on the SCMR problem is limited along two axes, *optimality* and *generality* (see Table 1 for a summary): (1) *optimality*: some techniques do not optimize execution time [62], or optimize routing with respect to a fixed, trivial mapping [8]; (2) *generality*: other techniques [31, 33, 67] do not account for routing T gates, a major source of two-qubit gates, and/or assume a sparse grid layout [8, 31, 33]. This means they cannot be directly applied to other architectures which pack qubits into a smaller grid area [44, 62]. Compact architectures are critical because they require fewer physical resources, but they make SCMR more challenging.

We present an *optimizing* mapping and routing algorithm that addresses the full *generality* of the SCMR problem, targeting circuits with T gates and architectures with arbitrary layouts.

The DASCOT approach. Obtaining optimal or near-optimal SCMR solutions is critical for both efficiency and correctness of quantum computation, but the problem is NP-complete [67]. We find that brute-force solutions based on, e.g., satisfiability solvers, do not scale to realistic problem sizes, as previously shown in the analogous mapping and routing problem for near-term noisy devices without error correction [42, 45, 60, 63].

We present a new algorithm called DASCOT (Dependency-Aware Surface Code Optimization) which mitigates the complexity of SCMR by breaking mapping and routing into independent subproblems, as illustrated in Fig. 2. The key insight underlying DASCOT is *dependency-awareness*: we exploit the dependency structure between circuit operations to carefully optimize the mapping and routing decisions.

Dependency-aware mapping. Our goal with mapping is to map circuit qubits in such a way that is less likely to cause conflicts in routing gates. To do so, we need to track which pairs of qubits interact in the circuit. While this idea has appeared in different guises in the qubit mapping literature,

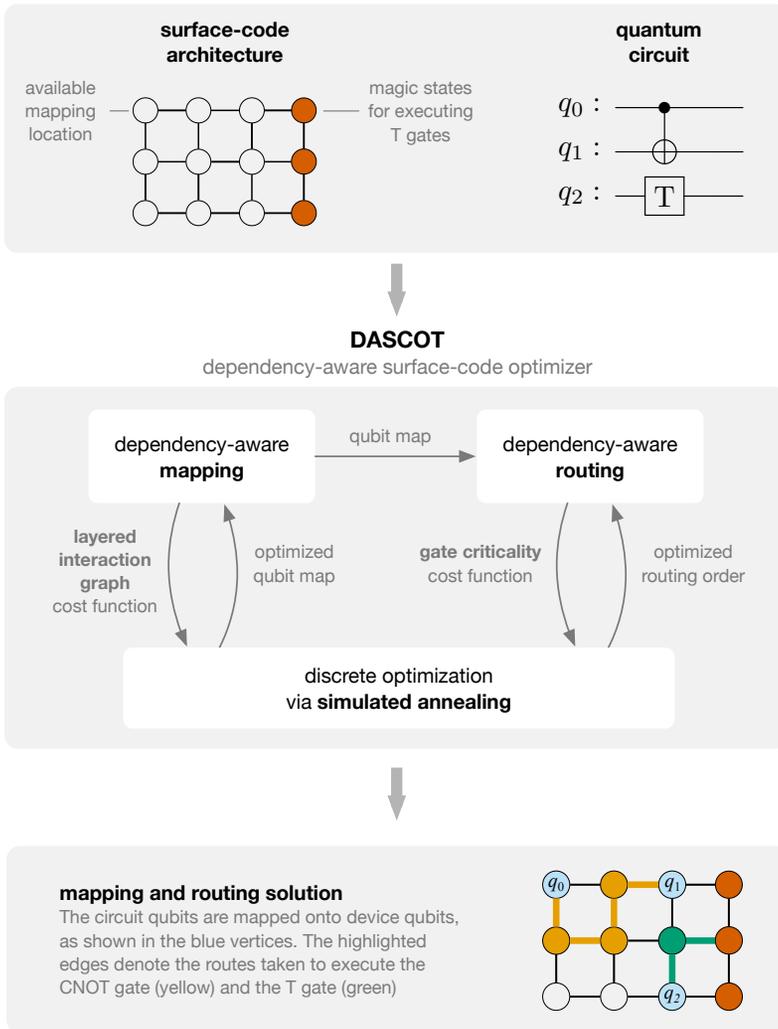


Fig. 2. Overview of our approach, DASCOT

existing solutions tend to be *flow-insensitive* in nature [31, 33], meaning that qubit interactions are modeled at a rough granularity that does not take into account gate ordering in the circuit.

We make the key insight that to build a good candidate mapping, we need finer-grained, flow-sensitive modeling of qubit interactions. For example, if one gate logically depends on another, then no valid routing solution will execute the two gates simultaneously, so there is no reason for the mapping stage to avoid crossing paths between them. By distinguishing between potential crossings which are likely to be problematic at the routing step and those which are not, we can better evaluate a candidate mapping. We capture this information in a new data structure we call a *layered interaction graph*, which is the basis of our *dependency-aware* mapping approach.

Dependency-aware routing. At the routing stage, our key insight is that the order in which gates are routed is crucial to the optimality of the scmr solution. Specifically, we formulate a discrete optimization problem over routing orders, with the goal of discovering a routing order that routes

gates that are *critical* in the dependency structure of the circuit. A critical gate has lots of other gates that consume its results later in the circuit, so routing it early unlocks deeper layers of the circuit for routing. Highly critical gates are prioritized to avoid a long “tail” of unexploited parallelism as the gates on the critical path are executed in sequence. We observe that greedily routing gates in order of criticality, as considered in prior work [33], does not maximize the *total* criticality of routed gates because a single gate can block the execution of multiple gates of slightly lower criticality. However, our new perspective on the routing problem as a search over the space of routing orders allows us to find the order which is best with respect to total criticality.

Reductions to discrete optimization. The mapping and routing subproblems each induce a discrete optimization problem with respect to a cost function. For mapping, we are searching for the mapping which minimizes the overlaps between gates in the same layer. For routing, we seek the gate ordering which maximizes the progress towards executing the entire circuit. In both cases, we solve these optimization problems with *simulated annealing* [35], a classic, easy-to-implement algorithm for solving hard, discrete optimization. Additionally, simulated annealing supports different tradeoffs between compute time and solution quality.

SAT-based optimal baseline. We present an optimal algorithm for the SCMR problem based on a novel SAT encoding of SCMR. Much like constraint-based approaches to mapping and routing for quantum computers without error correction [42, 53, 60], we construct a boolean formula that is satisfiable if and only if there exists an SCMR solution with a given execution time. Then, we find an optimal solution through a sequence of calls to a SAT solver, starting from a known lower bound and incrementing until the formula is satisfiable. Though the SAT-based approach does not scale to large-scale applications, it allows us to find exact solutions for small circuits and to precisely assess the optimality of DASCOT for these instances.

Evaluation. We present a thorough evaluation of DASCOT on a comprehensive suite of benchmarks consisting of over 200 circuits implementing key quantum algorithms and target two architectures representing different space–time trade-offs. Our results confirm the efficiency and effectiveness of DASCOT. For example: (1) in terms of solution quality, DASCOT matches or outperforms a state-of-the-art approach for SCMR without T gates, Autobraid [31], on 84% of benchmarks, with significant gains on *dense* circuits; including up to 31% cost improvement on Quantum Fourier Transform circuits; (2) in 69 circuits where we can compute an optimal solution using a satisfiability solver, DASCOT is within 25% of optimal for all but 5; (3) DASCOT can solve benchmarks with thousands of gates within minutes and can be tuned to different tradeoffs between compute time and optimality.

Contributions. Our contributions are the following:

- A formalization of the surface code mapping and routing problem (SCMR) capturing the combinatorial task of compiling circuits for surface code devices (Section 3). Compared to existing work, our formalization models the problem in more detail and with more flexibility, explicitly modeling magic states and allowing for arbitrary grids.
- A dependency-aware mapping algorithm using a new data structure called a *layered interaction graph* (Section 4.2). In comparison to prior work [31, 33], the layered nature of our interaction graphs captures finer-grained—*flow-sensitive*—interaction patterns between qubits, allowing us to choose better candidate mappings that maximize parallelism.
- We make the key observation is that routing order is crucial for the optimality of SCMR solutions. We therefore present a dependency-aware routing algorithm which searches over routing orders and chooses the next gates to route based on progress through the dependency structure of the circuit (Section 4.3). This is in contrast to greedy approaches [31] that choose

to route gates by order of their importance, disregarding the downstream effects of such actions, and end up with suboptimal global solutions.

- A SAT-based SCMR algorithm for finding optimal solutions on small circuits (Section 5). To our knowledge, this is the first optimal solver for the mapping and routing problem on a fault-tolerant architecture. The SAT-based algorithm allows us to quantify the distance from optimality for our approximate, simulated-annealing-based solutions.
- An empirical evaluation demonstrating the optimality and generality of DASCOT (Section 6).

2 Surface Code Compilation: A Primer

In this section, we describe the surface code mapping and routing problem, present illustrative examples, and provide a high-level overview of our approach.

Quantum error correction. Physical realizations of qubits are extremely delicate and subject to unintended changes of state, leading to computational errors. The leading approach to address errors is called a surface code [24]. A surface code encodes a logical qubit using a two-dimensional lattice of physical qubits; an example of a logical qubit encoded using a surface code is shown on the left of Fig. 3. Each physical qubit in the lattice is designated as either a data qubit (large circles in the figure) or a measurement qubit (small circles). Data qubits carry the state of the logical qubit, whereas measurement qubits are repeatedly measured to detect errors. There are two types of measurement qubits (indicated by color of the surrounding region in the figure) to detect both bit-flip and phase-flip errors. As depicted on the right of Fig. 3, a surface code quantum device consists of multiple surface code logical qubits encoded into one large lattice of physical qubits.

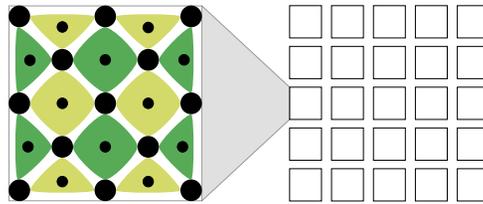


Fig. 3. Quantum device implementing a surface code

2.1 Surface Code Mapping and Routing by Example

We will now introduce the surface code mapping and routing problem by developing some examples. The grid from Fig. 3 will be our quantum computing architecture. The logical qubits in the grid may be used for storing qubits of a circuit or, as we will see shortly, for routing CNOT gates.

Suppose we wish to execute the circuit consisting of two parallel CNOT gates shown in the top left of Fig. 4. First, we need to choose a logical qubit to encode each of the four qubits in the circuit. This is called a *mapping*. One possible mapping is shown below the circuit. Then, we need to plan the execution of the gates. There are two main proposed implementations for executing CNOT gates on two logical qubits: defect braiding [24] and lattice surgery [23, 30]. Both require establishing a path between the qubits. We focus on lattice surgery because of recent evidence that it is a more resource-efficient paradigm [23], but our core approach can also be applied to architectures based on the defect-braiding CNOT gate.

To apply a lattice surgery CNOT, we need to find a path of logical qubits on the grid, called ancilla qubits, from the control qubit to the target qubit. The path must connect a horizontal boundary of the control (the top or bottom edge) to a vertical boundary of the target, making at least one “bend”. The process of reserving ancillae for gates is called *routing*. In our case, one CNOT gate requires

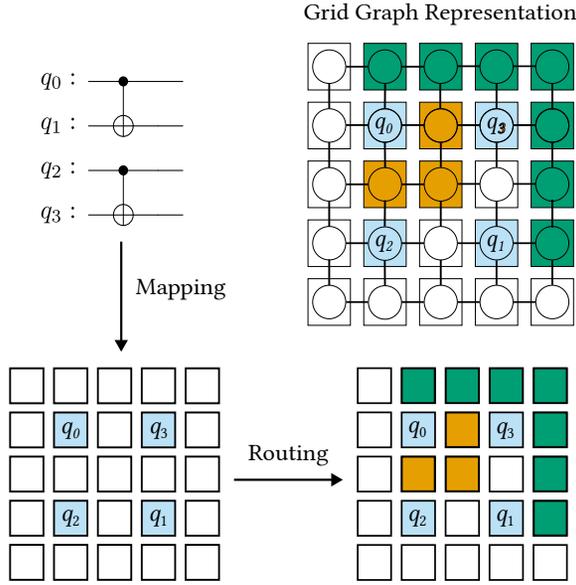


Fig. 4. Simple instance of scmr

an ancilla connection between q_0 and q_1 , while the other requires one between q_2 and q_3 . As long as these two connections do not overlap, the gates can be performed simultaneously in a single *time step*. A routing solution that meets these requirements is shown in the bottom right of Fig. 4. The two routes are indicated by colored squares, with the longer green route corresponding to the CNOT gate applied to q_0 and q_1 .

We can represent surface code mapping and routing with a grid graph as shown in the upper right of Fig. 4, which corresponds to the same instance and solution. Specifically, each vertex in the grid represents a logical qubit, and we include an edge between each adjacent pair of logical qubits, not including diagonals.

Optimal mapping & routing. The above example was fairly unconstrained. Any mapping leaving an ancilla space between each circuit qubit has a corresponding optimal routing solution. This is not the case in general. Consider the example in Fig. 5 wherein we are mapping the same circuit onto a 3×3 subset of our architecture. The depth of this circuit is 1, so it is theoretically executable in a single time step. However, with the choice of the mapping on the left labeled “Mapping 1”, there is no way to simultaneously execute the CNOT gates, as any paths from q_0 to q_1 and q_2 to q_3 will have to cross. Therefore, execution of the circuit will take two time steps. If instead we choose “Mapping 2”, we see that there is a routing solution that allows for simultaneous execution, with the routes indicated by colored paths.

Executing T gates. The τ gate cannot be applied directly to surface code logical qubits. Instead, executing a τ gate requires applying a lattice surgery CNOT between the input to the τ gate and a logical qubit prepared in a so-called “magic state” [12]. We make a standard assumption that magic state qubits are prepared in a separate region of the plane at a sufficiently high rate such that they are always available at dedicated storage locations [8]. The implication for our mapping and routing problem is that *each T gate must be routed like a CNOT gate with the target qubit chosen from a fixed set of vertices provided as part of the input.*

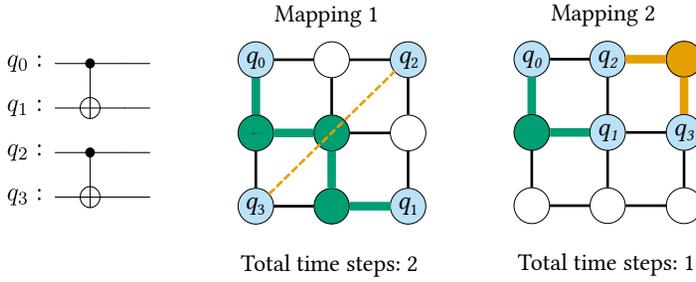


Fig. 5. Suboptimal mapping increases execution time

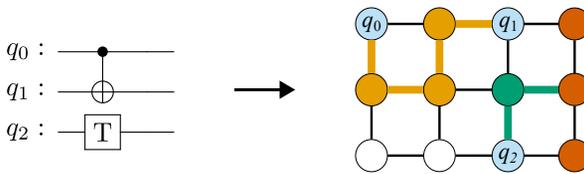


Fig. 6. Example with a T gate

We provide a simple example in Fig. 6. Since this circuit contains a T gate, we need to define magic state qubit locations. The 3×3 architecture has been extended with a column of magic state qubits along the right side, indicated with orange vertices. An optimal mapping and routing solution with one time step is shown. We have two simultaneous connections. One is between qubits q_0 and q_1 and corresponds to the CNOT, while the other is between q_2 and a magic state qubit, corresponding to the T gate.

2.2 An Overview of Our Approach

We now provide an overview of our approach (see Fig. 2).

Architectural flexibility. There is a fundamental trade-off between space and time in choosing a surface code architecture. A large grid with many available routing ancillae likely leads to shorter execution time, as many disjoint paths are available to route gates in parallel. A denser grid may save on qubits, but extend execution time. The appropriate balance of qubit footprint and execution time may vary depending on characteristics of the underlying hardware and the computation to be performed. Moreover, physical qubits are prone to defects [39, 41] which may force adapting the target architecture to omit certain regions of the grid.

We design our approach to support arbitrary architectures. For example, given a 9 qubit circuit, Fig. 7 shows two standard architectures proposed in the literature [30, 43, 62] which we call the Square Sparse and Compact architectures. These two architectures represent opposite ends of the space-time spectrum, with the Square Sparse architecture using many more qubits to allow routing more gates in parallel.

Dependency-aware mapping. To determine whether a mapping is appropriate for a circuit, we need to estimate the number of times where routing one gate might prevent the routing of another. For example, consider the circuit in Fig. 8, which is a simplified version of the 4 qubit Quantum Fourier Transform. The depth of this circuit is 5; equivalently, it can be partitioned into 5 layers of logically parallel gates (indicated with dashed lines).

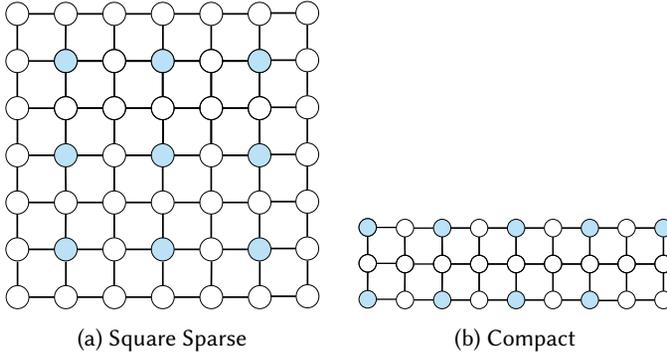


Fig. 7. Architectures with different space-time tradeoffs

Ideally, a mapping solution should enable routing the gates of each layer simultaneously. Our approach efficiently searches the space of mappings for one that minimizes contention between gates within the same layer using a simple data structure we call a *layered* interaction graph. Here, the only nontrivial layer is layer 3. To allow optimal routing of this layer, our mapping algorithm reserves disjoint regions (indicated with the dashed boxes) for the two gates, ensuring there exist nonoverlapping paths along which to route them. We call this a *dependency-aware* algorithm because it uses the dependency structure of the circuit to group gates into layers and establish which pairs of gates may be competing for routing resources.

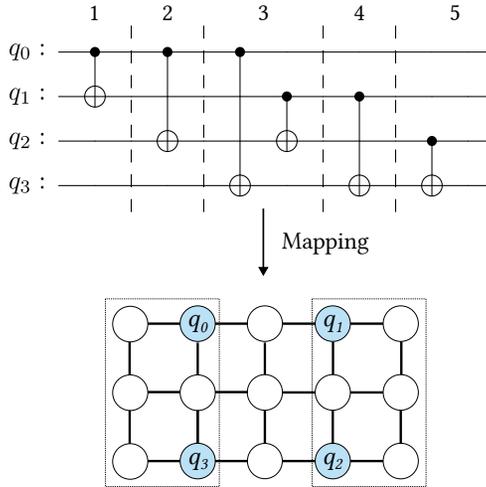


Fig. 8. Illustrating dependency-aware mapping

Dependency-aware routing. Once the qubit map is fixed, we proceed by iteratively routing as many logically parallel gates as possible. We can reduce the space of possible routing solutions by routing each gate in sequence, along the shortest available path. With this strategy, the order in which we iterate over the gates is important.

Consider Fig. 9 with our simple circuit with one τ gate and one CNOT. If we first route the CNOT gate along the shortest path, then the τ gate must wait for the next time step as there is no path of

available vertices between q_2 and a magic state qubit. On the other hand, if we start with the τ gate, then the CNOT can still be routed around the perimeter of the architecture, yielding a single time step.

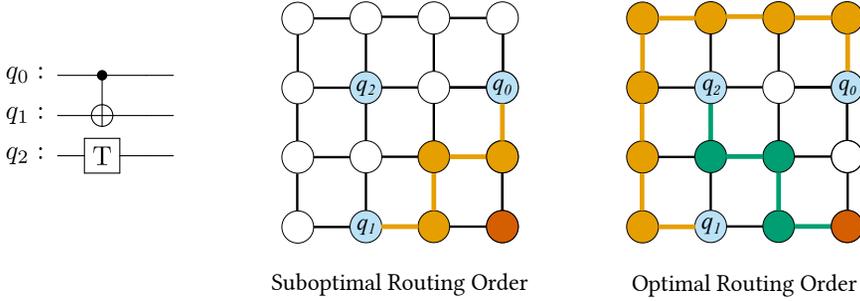


Fig. 9. Illustrating dependency-aware routing

The key question is therefore: *How do we choose the best routing order for a set of parallel gates?* We formulate this challenge as a discrete optimization problem over routing orders. We search for an order which allows the simultaneous routing of many gates to make progress towards execution of the complete circuit. Additionally, we prefer to route *critical* gates, those with long chains of dependent gates. Weighing gates by criticality prevents cascading effects from delaying a gate with many dependencies.

At the heart of dependency-aware mapping and routing is a minimization subroutine. Since our problem is discrete in nature and infeasible to solve optimally, we use *simulated annealing* to obtain solutions which approach the minimum value. Simulated annealing is a classical, iterative optimization algorithm that allows us to trade compute time for solution quality by tuning the number of iterations.

3 Surface Code Mapping & Routing

We now formalize surface code mapping and routing.

Architectures. Our abstraction for a fault-tolerant quantum computing architecture is a grid graph with some of the vertices reserved for magic state qubits and others available for storing circuit qubits and routing gates. We use $A = (V, E, MS)$ to denote an architecture where $G = (V, E)$ is a grid graph and $MS \subset V$ are vertices reserved for magic state qubits.

Formally, the $m \times n$ grid graph is defined over the vertex set $V = \{1, \dots, m\} \times \{1, \dots, n\}$. The graph includes an edge between any pairs of vertices (a, b) and (c, d) such that either $|a - c| = 1$ or $|b - d| = 1$. In the former case, we say that these vertices are *horizontal neighbors*. In the latter, we say that they are *vertical neighbors*.

Qubit maps. Given a circuit C acting on a set of qubits Q and an architecture $A = (V, E, MS)$, a *qubit map* is a function, denoted M , that uniquely maps each qubit in Q to one of the non-magic-state qubits on the device (i.e., vertices $V \setminus MS$).

Example 3.1. Fig. 8 shows a qubit map where four qubits are mapped to four of the 15 vertices in the architecture.

Dependency and routing. Let $C = \langle g_1, \dots, g_k \rangle$ denote a circuit represented as a sequence of gate applications. If two gates g_i and g_j act on a shared qubit and $i < j$, then we say g_j *depends* on g_i . For example, in the circuit in Fig. 8, the gate CNOT $q_1 q_2$ depends on the gate CNOT $q_0 q_1$. For routing,

we must assign gates to execution time steps in a way that respects dependencies in C . Also, it must be possible to execute all of the gates assigned to any particular time step simultaneously on the architecture. Gates can be executed simultaneously if there are vertex-disjoint paths, of the proper shape, along which to route them. We say that an assignment of gates to time steps and paths is a *valid gate route* if it meets these requirements. In the definition to come, $\langle V \rangle$ denotes the set of finite sequences of vertices $\langle v_1, \dots, v_m \rangle$ from a vertex set V .

Valid gate routes. Given an architecture $A = (V, E, MS)$, circuit C , and qubit map M , a *valid gate route* is a natural number t representing a number of time steps, and a pair of functions $R_{\text{SPACE}} : C \rightarrow \langle V \rangle$ and $R_{\text{TIME}} : C \rightarrow \{1, \dots, t\}$ that meets the following criteria:

- **Data Preservation:** No vertex representing a circuit qubit or magic state qubit is used for routing by R_{SPACE} .
- **CNOT Routing:** If $g = \text{CNOT } q_i q_j$, then $R_{\text{SPACE}}(g)$ is a path from a vertical neighbor of $M(q_i)$ to a horizontal neighbor of $M(q_j)$.
- **T Routing:** If $g = \text{T } q_i$, then $R_{\text{SPACE}}(g)$ is a path from a vertical neighbor of $M(q_i)$ to a horizontal neighbor of a magic state.
- **Gate Order:** If g_j depends on g_i , $R_{\text{TIME}}(g_i) < R_{\text{TIME}}(g_j)$.
- **Disjoint Paths:** If $R_{\text{TIME}}(g_i) = R_{\text{TIME}}(g_j)$, then the paths $R_{\text{SPACE}}(g_i)$ and $R_{\text{SPACE}}(g_j)$ share no vertices.

Example 3.2. Consider Fig. 9. Let g_1 be the gate $\text{CNOT } q_0 q_1$ and g_2 be the gate $\text{T } q_2$. With the optimal routing order on the right, we show a valid gate route with $t = 1$ time steps, so $R_{\text{TIME}}(g_1) = R_{\text{TIME}}(g_2) = 1$. The path $R_{\text{SPACE}}(g_1)$ is the longer yellow path, while $R_{\text{SPACE}}(g_2)$ is the shorter green path.

With this terminology in place, the surface code mapping and routing problem (SCMR) can be concisely defined as the task of finding a qubit map and gate route pair with the minimal number of time steps.

The SCMR problem: Given an architecture A and circuit C , find a qubit map M and a valid gate route $(t, R_{\text{TIME}}, R_{\text{SPACE}})$ such that t is minimized.

4 The DASCOT Algorithm

The SCMR problem is NP-complete [67], so exact optimization is infeasible for large instances. To simplify the problem, we decouple the mapping and routing and solve each separately. We present novel algorithms for mapping and routing that exploit circuit dependency structure and use simulated annealing to efficiently search the space of solutions. We start with a short primer on simulated annealing.

4.1 Simulated Annealing Primer

In both mapping and routing, we define a *cost function* f and search for qubit maps and routing orders (respectively) with lower cost. Therefore DASCOT must efficiently solve problems of the form $\min_{s \in \mathcal{S}} f(s)$, where \mathcal{S} is a large discrete set.

For this purpose, we employ *simulated annealing*, a standard and simple algorithm well-suited to approximating the global optimum for discrete problems with large search spaces and many local minima. Simulated annealing begins by choosing a random candidate solution s_{curr} in \mathcal{S} and an initial temperature τ_0 . Then, it iteratively applies the following probabilistic local search step, repeating until the current temperature τ reaches some threshold τ_f .

- Uniformly sample a random neighbor s_{next} of s_{curr}

- If $f(s_{next}) < f(s_{curr})$, set s_{curr} to s_{next}
- Otherwise, do so with probability that is a function of the current temperature and cost increase
- Reduce the temperature by a cooling rate r

The non-zero probability of transitioning to a worse solution allows escape from local minima. We choose the standard acceptance probability $\exp\left\{-\frac{f(s_{new})-f(s_{curr})}{\tau}\right\}$. Thus, transitions to a worse solution are more likely early in the optimization, when the temperature is relatively high.

To instantiate simulated annealing, we need to define the cost function f , solution space \mathcal{S} , and $N(s)$, the set of neighbors of a solution s . The algorithm also has hyperparameters: the initial temperature τ_0 , termination temperature τ_f , and cooling rate r ; we describe our choices for these in Section 6.

4.2 Dependency-Aware Mapping

In the mapping stage, our goal is to find a qubit map which is likely to admit a valid gate route with few steps without explicitly computing the full routing. We present an approach based on a novel data structure called a *layered interaction graph*. Whereas prior work takes a “flattened” view of qubit interactions with no consideration of the ordering of gates, layered interaction graphs guide our search for a mapping using the dependency structure of a circuit.

Interaction graphs. A natural starting point for selecting a qubit map is to analyze which pairs of qubits interact by constructing an *interaction graph*. [31, 33]. An interaction graph for a circuit C includes a vertex for each qubit and an undirected edge (q_i, q_j) for each pair such that $\text{CNOT } q_i q_j$ is a gate in C . To capture interactions with magic states, we also include an additional vertex v_t and add an edge (q_i, v_t) whenever there is a τ gate on the qubit q_i .

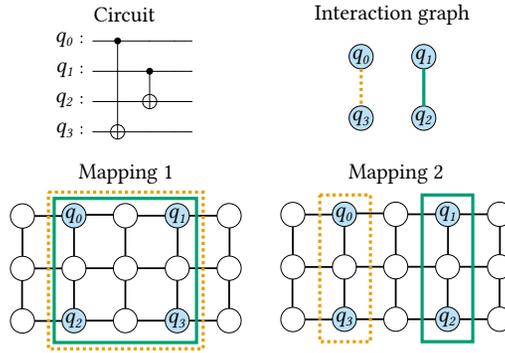


Fig. 10. Interaction graphs and bounding boxes

Example 4.1. Fig. 10 shows a circuit of two parallel CNOT gates. The interaction graph for this circuit (upper right) consists of two edges, one between q_0 and q_3 and another between q_1 and q_2 .

Gate conflicts. We can evaluate the suitability of a map for a circuit using its interaction graph. For a given mapping, we can find the *bounding box* of the interaction graph edges, following prior work [31]. The bounding box of an edge is defined as the smallest rectangular region on the architecture which contains both of the interacting qubits at its endpoints. We say two gates *conflict* when their bounding boxes overlap because the routing of one gate may block the routing of the other.

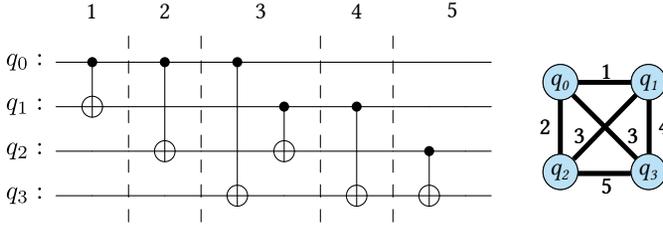


Fig. 11. A circuit and its layered interaction graph

Example 4.2. In the second row of Fig. 10, we compare two qubit maps using the interaction graph. The bounding boxes of the two interaction graph edges with respect to each map is shown in the same color and style. Following our heuristic, we prefer Mapping 2, where the gates do not conflict, to Mapping 1, where they do. This is in fact the correct choice because there is a valid gate route with one step for Mapping 2 but not for Mapping 1.

Layered interaction graphs. However, the problem with this strategy is that it discards important information for routing: the dependency between gates. To address this limitation, we present *layered* interaction graphs. We first partition a circuit into a sequence of *layers* such that if a gate g depends upon another g' , then g must be assigned to a later layer than g' . Then, we build an interaction graph where each edge is labeled with the layer that the corresponding gate is assigned to. The resulting graph with edge labels is a *layered* interaction graph.

Example 4.3. A circuit based on the Quantum Fourier Transform and its layered interaction graph are shown in Fig. 11. Notice that this circuit has 5 layers, and layer 3 is exactly the circuit from our previous example.

To illustrate the advantage of layered interaction graphs, in Fig. 12, we return to the same two qubit maps and evaluate them with respect to the circuit and corresponding layered interaction graph from Example 4.3.

The only parallel gates in this circuit are the two gates in layer 3, so Mapping 2 remains a better choice because it enables routing these gates in the same time step. Thus, an ideal cost function should assign lower cost to Mapping 2 than Mapping 1. Counting conflicts with respect to the standard interaction graph instead assigns the same cost to these maps because both induce one conflict, indicated by the yellow dashed edges. However, in the *layered* interaction graph, we can distinguish between these cases using the edge labels. If we only consider conflicts between edges with the same label, we see one conflict under Mapping 1 in layer 3 (as in Fig. 10), and no conflicts under Mapping 2 (the edge labels 2 and 4 do not match), recovering the information that Mapping 2 is the better choice.

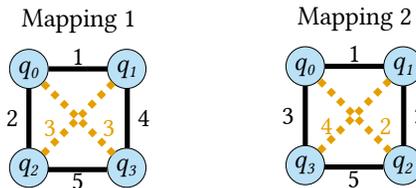


Fig. 12. Qubit maps and layered interaction graph

This is the key insight behind our algorithm, shown in Alg. 1. We construct a layered interaction graph (lines 2-8), then minimize *conflicts* with respect to the layered interaction graph. Two edges are said to *conflict* under a map M if their bounding boxes overlap and they have the same layer label. The bounding box of an edge associated with a gate $\text{CNOT } q_i q_j$ is the bounding box of the architecture vertices $M(q_i)$ and $M(q_j)$. The bounding box of an edge associated with a gate τq_i is the bounding box of the architecture vertices $M(q_i)$ and the nearest magic state vertex.

Algorithm 1 Dependency-Aware Mapping

```

1: procedure DASCOT-MAP(arch  $A$ , circuit  $C$ , qubits  $Q$ )
2:   Partition  $C$  into layers  $\mathcal{L} = L_1, \dots, L_d$ .
3:   Let  $I$  be the empty graph over the vertices  $Q \cup \{v_t\}$ 
4:   for layer  $L_k$  in  $\mathcal{L}$  do
5:     for each  $\text{CNOT } q_i q_j$  in  $L_k$  do
6:       Add edge  $(q_i, q_j)$  to  $I$  with label  $k$ 
7:     for each  $\tau q_i$  in  $L_k$  do
8:       Add edge  $(q_i, v_t)$  to  $I$  with label  $k$ 
9:   return  $\text{argmin}_M \text{CONFLICTS}(I, M)$ 

```

Solving the minimization objective. The minimization problem (blue in Alg. 1) can be solved via simulated annealing: (1) the search space \mathcal{S} is the set of possible qubit maps, (2) the neighbor set of $N(M)$ of a map M is those obtainable by swapping the locations of one pair of circuit qubits, and (3) the cost function f is the number of conflicts between edges with the same label in the interaction graph.

4.3 Dependency-Aware Routing

At the routing stage, we take a circuit, architecture, and qubit map as input and return a sequence of time steps. As formalized in the definition of a valid gate route in Section 3, each time step k consists of (1) a set of parallel gates, corresponding to the gates g for which we set $R_{\text{TIME}}(g) = k$, and (2) disjoint paths along which to route them, defining the function $R_{\text{SPACE}} : C \rightarrow \langle V \rangle$ on the same set of gates. The goal is to find the sequence with the fewest total time steps. The distinguishing feature of our approach is that we define and solve a discrete optimization problem over the space of candidates for the each time step, focusing on gates with high *criticality*.

We begin by describing our algorithm for constructing a single time step for a given layer of parallel gates, which we will iteratively apply to generate a full routing solution. To limit the combinatorial explosion in possible routes, we consider gates sequentially, routing each along the shortest available path, if one exists (otherwise the gate is delayed to a future time step). Alg. 2 shows this subroutine for a particular sequence σ of gates.

With this linear approach, the order in which we iterate over gates affects the resulting time step because a gate earlier in the order may occupy vertices which prevent the routing of a later gate. Therefore, our goal is to find the routing order σ that results in the best time step.

Shortest-first order. One approach to a routing order is to choose the next gate to be the one with the shortest available routing path. Intuitively, this order is often effective because shortest paths occupy fewer vertices and are thus less likely to block other paths. In fact, the shortest-first order has a guaranteed approximation ratio derived from the literature on the *node-disjoint paths* problem [8, 37]. It results in routing $O(\text{OPT}/\sqrt{n})$ gates where OPT is the optimal number of gates which can be routed in the next time step and n is the number of vertices in the architecture.

Algorithm 2 Single Step Routing

```

1: procedure ROUTE-SEQ(gate sequence  $\sigma$ , arch  $A$ ,  $M$ )
2:   initialize  $R'_{\text{SPACE}}$  to  $\emptyset$ 
3:   for gate  $g$  in  $\sigma$  do
4:     Let  $p$  be a shortest path for routing  $g$  in  $A$ 
5:      $R'_{\text{SPACE}} \leftarrow R'_{\text{SPACE}} \cup \{(g, p)\}$  ▷ if some path exists
6:     Remove all vertices in  $p$  from  $A$ 
7:   return  $R'_{\text{SPACE}}$ 

```

Routing order search. While ordering the gates by path length is a sound principle, there are numerous cases in practice where it produces suboptimal results.

Example 4.4. Consider Fig. 9. The shortest paths available to execute the two gates are equal. With shortest-first sorting, these two orders are equivalent. However, only one order allows routing both gates in the same time step.

To prevent cases like this, where a fixed heuristic cannot identify an optimal routing order, we search for the best routing order at each iteration, generating several candidate time steps and picking the best. To define the “best” choice among candidate time steps, we need to estimate the degree to which a time step will lead to a high-quality overall solution, with few total time steps.

Dependency-aware cost. We use a dependency-aware metric, where the cost reduction associated with routing a gate g is equal to its *criticality*—we want to prioritize routing important gates first.

Criticality: Let C be a quantum circuit and g be a gate in C . The *criticality* of g , denoted $\text{CRITICALITY}(g)$ is the depth of the subcircuit of C consisting of gates which depend on g , including g itself.

For example, in the circuit in Fig. 6, both gates have a criticality of 1, while the gate CNOT $q_0 q_3$ in Fig. 8 has a criticality of 3. Intuitively, routing critical gates unlocks deeper layers of the circuit for routing. Gates with high criticality should be prioritized to avoid a long “tail” of underutilized time steps as the gates on the critical path are executed in sequence.

Thus, we search for the routing order such that the resulting time step $s = \text{ROUTE-SEQ}(\sigma, A, M)$ is minimal with respect to the dependency-aware cost function below (highlighted in blue in Alg. 3)

$$f(s) = \sum_{(g,p) \in s} -\text{CRITICALITY}(g)$$

Note that this cost cannot be minimized greedily. The critical-first order, routing gates in order of criticality, is not necessarily optimal.

Example 4.5. Fig. 13 depicts the front layer of a circuit with three parallel gates. Suppose the CNOT gate has a criticality of 3 and the two τ gates have criticality 2. By first routing the most critical gate, the CNOT, we block the routing of the two τ gates. Therefore the critical-first routing order has a cost of -3. On the other hand, if we first consider the τ gates, both can be routed in the same time step, yielding a cost of -4. Thus, the greedy approach is not optimal here.

Our routing algorithm is Alg. 3. It determines the front layer of the circuit, applies simulated annealing to find the best choice for the next time step, and repeats until all gates have been routed, returning a sequence of time steps. We instantiate simulated annealing as follows: (1) the search space \mathcal{S} is the set of orderings of the layer, (2) the neighbor set of $N(\sigma)$ of an ordering σ is those

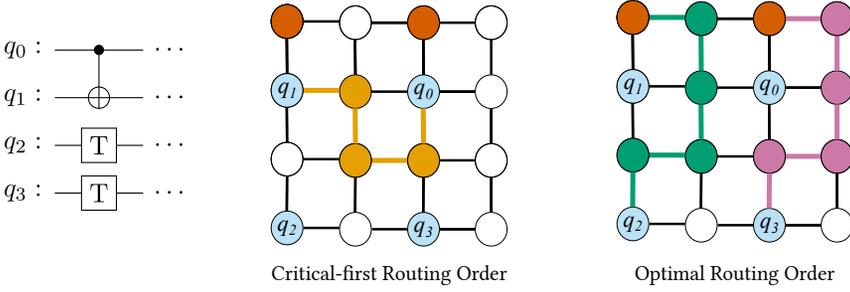


Fig. 13. A case where the critical-first order is suboptimal

Algorithm 3 Dependency-Aware Routing

```

1: procedure DASCOT-ROUTE(arch  $A$ , circuit  $C$ , map  $M$ )
2:   initialize  $R_{\text{TIME}}$  and  $R_{\text{SPACE}}$  to  $\emptyset$ 
3:   initialize number of steps  $t$  to 0
4:   while  $C \neq \emptyset$  do ▷ note we remove gates from  $C$ 
5:     Let  $L_1$  be the front layer of gates in  $C$ 
6:     Let  $\Sigma$  be the set of all possible orderings of  $L_1$ 
7:      $\sigma^* \leftarrow \operatorname{argmin}_{\sigma \in \Sigma} f(\text{ROUTE-SEQ}(\sigma, A, M))$  ▷ simulated annealing
8:      $R_{\text{NEXT}} \leftarrow \text{ROUTE-SEQ}(\sigma^*, A, M)$ 
9:      $t \leftarrow t + 1$ 
10:    for each (gate, path) pair  $(g, p)$  in  $R_{\text{NEXT}}$  do
11:       $R_{\text{TIME}}(g) \leftarrow t$ 
12:       $R_{\text{SPACE}}(g) \leftarrow p$ 
13:      Remove  $g$  from  $C$ 
14:  return  $R_{\text{TIME}}, R_{\text{SPACE}}$ 
    
```

obtainable by swapping the position of one pair of gates, and (3) the cost function is given by $f(\text{ROUTE-SEQ}(\sigma, A, M))$.

5 A SAT-Based Optimal Baseline

To evaluate the quality of solutions produced by DASCOT, we developed a SAT-based optimal solver for the SCMR problem. Though the SAT approach does not scale to large circuits, it provides a “ground truth” for small circuits. Like approaches for NISQ computers [42, 53, 60], the optimal solver encodes the existence of a solution with k time steps as a query to a boolean satisfiability solver, incrementing k if the result is “unsatisfiable”, and iterating until a solution is found. In the rest of this section, we describe the encoding for a fixed value of k .

Throughout we will fix an architecture $A = (V, E, MS)$, a circuit $C = \langle g_1, \dots, g_k \rangle$ acting on qubits in a set Q and a time limit t_s . Our encoding includes three types of boolean variables: the map variables represent the qubit mapping, the exec variables represent the assignment of gates to time steps (the function R_{TIME}) and the path variables represent the path used to route gates (the function R_{SPACE}). The formula over these variables that we produce will enforce that the qubit map encoded by map and the routing encoded by exec and path are a valid solution as defined in Section 3. In describing these constraints, we use cardinality constraints which enforce that *at most one* (AMO) or *exactly one* (EO) of a set of variables is set to true. The efficient translation of cardinality constraints

into conjunctive normal form is a well-studied area [4, 6, 25, 49, 56] so we leave the encoding unspecified for simplicity.

Variables. The boolean variables that appear in our encoding are listed below along with their intended semantics.

- $\text{map}(q, v)$: the qubit q in the circuit is mapped to the vertex v .
- $\text{exec}(g, t)$: the gate g is executed in time step t .
- $\text{path}(u, v, g, t)$: the edge (u, v) is used in the path for gate g at time step t .

5.1 Constraints

Maps are injective functions. To enforce that a satisfying assignment corresponds to a valid qubit map, we require that for each q there is *exactly one* v where $\text{map}(q, v)$ is set to true (functional consistency) and that for each v there is *at most one* q where $\text{map}(q, v)$ is set to true (injectivity). These constraints are shown below.

$$\text{MAP-VALID} \triangleq \bigwedge_{q \in Q} \text{EO}(\{\text{map}(q, v) : v \in V \setminus MS\}) \wedge \bigwedge_{v \in V \setminus MS} \text{AMO}(\{\text{map}(q, v) : q \in Q\})$$

Paths preserve data. We cannot overwrite circuit qubits or magic state qubits in the routing process. In terms of path and map variables, this means that we cannot set both $\text{path}(u, v, g, t)$ and $\text{path}(v, w, g, t)$ for any u, w, g , and t if we have set $\text{map}(q, v)$ for some qubit q or if $v \in MS$.

$$\begin{aligned} \text{CIRC-SAFE} &\triangleq \bigwedge_{\substack{u, v, w \in V \\ t \leq t_s \\ q \in Q, g \in C}} \neg \text{map}(q, v) \vee \neg \text{path}(u, v, g, t) \vee \neg \text{path}(v, w, g, t) \\ \text{MS-SAFE} &\triangleq \bigwedge_{\substack{v \in MS, u, w \in V \\ t \leq t_s \\ g \in C}} \neg \text{path}(u, v, g, t) \vee \neg \text{path}(v, w, g, t) \\ \text{DATA-SAFE} &\triangleq \text{CIRC-SAFE} \wedge \text{MS-SAFE} \end{aligned}$$

All gates are executed in logical order. To ensure that all gates are executed in a valid order, we enforce that for each g , there is exactly one t such that $\text{exec}(g, t)$ is set to true and that if g_j depends on g_i , (which we denote with $g_i \prec g_j$), then $\text{exec}(g_i, t)$ implies $\text{exec}(g_j, t')$ for some $t' > t$. This is represented by the constraints:

$$\text{GATES-ORDERED} \triangleq \bigwedge_{g \in C} \text{EO}(\{\text{exec}(g, t) : t \leq t_s\}) \wedge \bigwedge_{\substack{g \prec g' \\ t' \leq t}} \neg \text{exec}(g, t) \vee \neg \text{exec}(g', t')$$

Paths are disjoint. Though the edges in the underlying graph are undirected, it is useful for our formulation to assign a direction to the paths, so we assign the direction of the edge as from u to v if $\text{path}(u, v, g, t)$ is set to true. We require that our simultaneous paths are *vertex disjoint*, meaning each vertex can have at most one incoming edge and at most one outgoing edge. Therefore, for each time step t and vertex u , we include the constraint that there is at most one pair v, g such that

$\text{path}(u, v, g, t)$ is set to true and at most one pair v', g' such that $\text{path}(v', u, g', t)$ is set to true:

$$\text{DISJOINT} \triangleq \bigwedge_{\substack{t \leq t_s \\ u \in V}} \text{AMO}(\{\text{path}(u, v, g, t) : v \in V, g \in C\}) \wedge \text{AMO}(\{\text{path}(v', u, g', t) : v' \in V, g' \in C\})$$

Paths connect CNOT pairs. The most complicated type of constraints are those that enforce that the path variables encode a valid path between the relevant vertices. We start with the case of CNOT gates. The constraints effectively construct a path inductively. Let $g = \text{CNOT } q_i q_j$ be a CNOT gate. The path corresponding to g must begin at a vertical neighbor of the control qubit q_i and terminate at a horizontal neighbor of the target qubit q_j . This leads to the “base case” constraint for each end of the path. We state them below using $VN(v)$ (and $HN(v)$) as an abbreviation for the set of up to two vertical (and horizontal) neighbors of a vertex v :

$$\begin{aligned} \text{CNOT-START} &\triangleq \bigwedge_{\substack{v \in V \\ g(q_i, q_j) \in C \\ t \leq t_s}} \left(\neg \text{map}(q_i, v) \vee \neg \text{exec}(g, t) \vee \bigvee_{u \in VN(v)} \text{path}(v, u, g, t) \right) \\ \text{CNOT-REACH-TARGET} &\triangleq \bigwedge_{\substack{v \in V \\ g(q_i, q_j) \in C \\ t \leq t_s}} \left(\neg \text{map}(q_j, v) \vee \neg \text{exec}(g, t) \vee \bigvee_{u \in HN(v)} \text{path}(u, v, g, t) \right) \end{aligned}$$

The inductive constraint is that for each edge (u, v) on the path, either u is a valid starting point for the path, or an internal vertex on the path. In the latter case, it must be incident to another edge to follow back towards the starting point. In our variables, we can express the two options as $\text{path}(u, v, g, t)$ implies either $\text{path}(w, u, g, t)$ for some distinct vertex w adjacent to u or $\text{map}(q_i, u)$.

$$\begin{aligned} \text{CNOT-INDUCTIVE} &\triangleq \bigwedge_{\substack{g(q_i, q_j) \in C \\ t \leq t_s \\ (u, v) \in E}} \left(\neg \text{path}(u, v, g, t) \vee \bigvee_{\substack{(w, u) \in E \\ w \neq v}} \text{path}(w, u, g, t) \vee \text{map}(q_i, u) \right) \\ \text{CNOT-ROUTED} &\triangleq \text{CNOT-START} \wedge \text{CNOT-REACH-TARGET} \wedge \text{CNOT-INDUCTIVE} \end{aligned}$$

Paths connect τ targets to magic states. Let $g = \tau q_j$ be a τ gate. We also need to enforce that the path variables assign an appropriate path between q_j and a magic state qubit. This type of constraint is similar to the previous. The only difference is in the definition of valid end points which are now horizontal neighbors of magic state qubits.

$$\begin{aligned} \text{T-REACH-TARGET} &\triangleq \bigwedge_{\substack{g(q) \in C \\ t \leq t_s}} (\text{EO}(\{\text{path}(u, v, g, t) : v \in MS, u \in HN(v)\})) \\ \text{T-ROUTED} &\triangleq \text{T-START} \wedge \text{T-REACH-TARGET} \wedge \text{T-INDUCTIVE} \end{aligned}$$

The formula corresponding to an SCMR instance is given by the conjunction of these constraints:

$$\varphi(A, C, t_s) \triangleq \text{MAP-VALID} \wedge \text{GATES-ORDERED} \wedge \text{DATA-SAFE} \wedge \text{DISJOINT} \wedge \text{CNOT-ROUTED} \wedge \text{T-ROUTED}$$

THEOREM 5.1. *The formula $\varphi(A, C, t_s)$ is satisfiable if and only if there is a map M and corresponding valid gate route $(t_s, R_{\text{TIME}}, R_{\text{SPACE}})$ for the SCMR problem given by A and C , and there is an explicit translation from SCMR solutions M , $(t_s, R_{\text{TIME}}, R_{\text{SPACE}})$ to models of $\varphi(A, C, t_s)$ and vice-versa.*

6 Implementation and Evaluation

We now present our evaluation of DASCOT.

Benchmarks. Our evaluation is performed on a benchmark suite of 232 application circuits. Our suite includes the entire set collected by Zulehner et al. [68]. This existing set consists of circuits derived from the RevLib suite [64] and programs written in the Quipper [28] and ScaffoldCC [34] quantum programming languages. We extended it with implementations of major quantum algorithms: Shor’s Algorithm [54], the Quantum Fourier Transform [18], Bernstein-Vazirani [7], QAOA [22], and Grover’s Algorithm [29].

Architectures. We target the two architectures Square Sparse and Compact from Fig. 7. The Square Sparse architecture represents a case with a large space footprint studied in prior work on mapping and routing [31, 62]. It is the smallest square grid which includes enough space to surround each circuit qubit with routing ancillae on all sides. For a circuit with n qubits, this results in a side length of $2\lceil\sqrt{n}\rceil + 1$. The Compact architecture, in contrast, is designed to minimize the required device qubits. This architecture is nearly linear, with three rows, and enough columns to include an ancilla qubit between each circuit qubit ($n - 1$ columns for even n). In both cases, we assume this mapping region is surrounded on all sides by magic state qubits.

Experimental setup. All runs were allotted a 1hr timeout on one core of an AMD EPYC™ 7763 2.45 GHz Processor and 32GB of RAM accessed via a distributed research cluster. Since DASCOT is a randomized algorithm, we collected 20 trials and plot the mean solution and a 95% confidence interval. We choose the simulated annealing parameters $\tau_0 = 100$, $\tau_f = 0.1/d$, $r = 0.1/d$ (where d is the depth of the circuit) for mapping and $\tau_0 = 10$, $\tau_f = 0.1$, $r = 0.1$ for routing. The constants were determined empirically with a grid search over the ranges $[1, 10^3]$ for τ_0 and $[10^{-3}, 1]$ for r and τ_f , using randomly generated circuits to avoid benchmark overfitting.

Research questions. Our experimental evaluation is designed to answer the following questions:

- RQ1: How does our approach compare to prior work on surface code mapping and routing?
- RQ2: What are the gains from our mapping and routing algorithms over greedy baselines?
- RQ3: How close to optimal are our solutions?
- RQ4: What is the tradeoff between optimality and scalability?

(RQ1) Comparison to Prior Work

To answer RQ1, we primarily compare DASCOT to Autobraid [31], a state-of-the-art algorithm for a closely related problem (though other tools are represented in addressing RQ2). Autobraid targets the defect-braiding CNOT gate and assumes a “sparse” architecture, like Square Sparse, where logical qubits are surrounded on all sides by routing ancillae. For a direct comparison, we implemented a version of Autobraid with two minor modifications: (i) constraints which require paths to include a “bend” as needed for lattice surgery and (ii) support for more general architectures. Moreover, *Autobraid does not handle routing of τ gates*. Therefore, for the purpose of this comparison only, we do not include them as part of the problem instance, as a special case of our SCMR problem. This matches the experimental setup originally used to evaluate Autobraid [31].

Overall results. We begin by analyzing the quality of solutions produced by Autobraid and DASCOT across the entire benchmark suite. In Fig. 14, we count the number of benchmarks where one algorithm outperforms the other in terms of solution quality by producing a *mean* solution with fewer time steps. The “Match” category represents circuits for which the two algorithms reached solutions with the same number of time steps.

On the Square Sparse architecture, the largest category by far is “Match” where the two algorithms reach equivalent solutions, containing 83% of the benchmark suite. In comparison, the “DASCOT

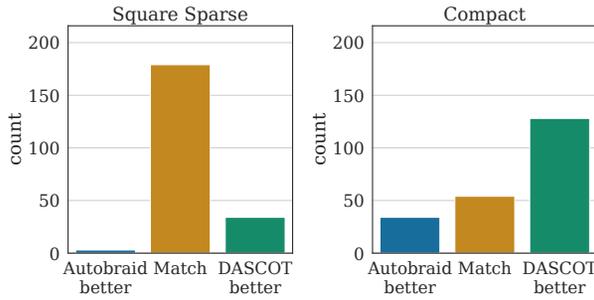


Fig. 14. Summary Comparison with Autobraid

“better” category contains 16% of circuits and “Autobraid better” category contains 1% (just three circuits). Most circuits are in the “Match” category because, with the plentiful routing resources of this architecture, both algorithms often achieve the theoretical lower-bound given by the depth of the circuit. On the other hand, the constraints of the Compact architecture do reveal a separation between the two algorithms. DASCOT outperforms or matches Autobraid on 84% of the benchmarks and strictly outperforms Autobraid on 59% of the benchmarks.

QFT circuits. Quantum Fourier Transform circuits are an important application which are particularly challenging to map and route because of their dense interaction graphs with many gates per layer. However, DASCOT is well-suited to the complexity of QFT circuits. Fig. 15 compares the *cost ratios* of the solutions provided by DASCOT and Autobraid for these hard instances. The cost ratio is defined as the quantity:

$$\frac{\text{\# of time steps in solution}}{\text{circuit depth}}$$

Since the depth of the circuit is a theoretical lower bound on the number of time steps in a solution, the best cost ratio is 1. DASCOT achieves lower mean cost ratios on than Autobraid up to about a 31% relative improvement on both architectures for the 100 qubit circuit.

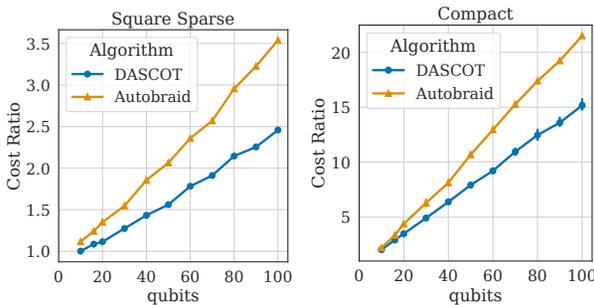


Fig. 15. Cost ratios on QFT Circuits

Synthetic dense circuits. We also generated a set of synthetic circuits to isolate the effect of density. To construct synthetic circuits, we iteratively applied a random layer of gates. For the “high-density” circuits, a gate is applied to each qubit in each layer. For the “low-density” circuits, a gate is only applied to 1/10 of the qubits in each layer. We generated circuits with up to 100 qubits and up to 100 layers.

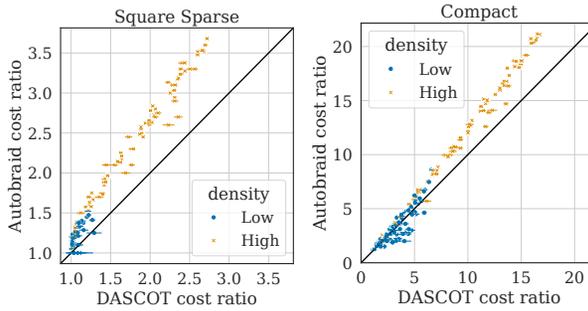


Fig. 16. Cost ratios on random circuits of varying density

Fig. 16 compares the cost ratios obtained by DASCOT and Autobraid for synthetic circuits. Each point represents a circuit. Points above the black line $y = x$ are those where DASCOT produces a better solution than Autobraid. For both algorithms, low-density circuits lead to lower cost ratios than high-density. This is expected because there is a high probability of conflict among the many pairs of simultaneous paths in the maximally dense layers. However, DASCOT is more robust to this complexity, outperforming Autobraid on almost all of the dense circuits across both architectures. We see a mean reduction in cost ratio of 22% on the Square Sparse architecture and 13% on the Compact architecture.

RQ1 Summary: DASCOT and Autobraid are both close to optimal on the Square Sparse architecture, while DASCOT outperforms the baseline on the Compact architecture, producing a strictly better solution for 59% of circuits and matching the Autobraid solution in a further 25%. We see especially significant gains from our approach on dense circuits, with up to a 31% improvement on QFT circuits.

(RQ2) Ablation Study

Next, we perform an ablation study to isolate the performance contributions of our mapping and routing algorithms. We focus on the Compact architecture and circuits with 8 or more qubits, since these are more difficult instances where a naive approach is unlikely to perform well.

First, to assess our mapping algorithm, we compare against a version of our algorithm called DASCOT-RANDMAP which chooses a qubit map at random, then applies the DASCOT routing algorithm (Alg. 3). Choosing a random map is equivalent to setting the initial temperature of the simulated annealing search equal to the final temperature, performing zero iterations of the search, or not applying a mapping optimization pass at all as in the lattice surgery compiler (LSC) [62].

In Fig. 17(left), we compare DASCOT to DASCOT-RANDMAP on our application circuits in terms of cost ratio. For the majority of circuits (60%), DASCOT produces strictly better results. For the remaining circuits, the search at the mapping phase does not improve the cost ratio, but there are no examples where it leads to a significantly worse cost ratio.

We see more uniform cost reduction for the routing phase. We analogously define the DASCOT-RANDROUTE algorithm which first applies the DASCOT mapping algorithm (Alg. 1). Then, for routing, it replaces the minimization procedure with a random order selection, once again corresponding to zero iterations of simulated annealing and the procedure applied by LSC [62]. In Fig. 17(middle), we compare DASCOT to DASCOT-RANDROUTE. The results indicate that a dependency-aware analysis of routing order leads to significantly stronger solutions than committing to a single random order at each step. We see a strict improvement in mean cost ratio on 90% of application circuits from

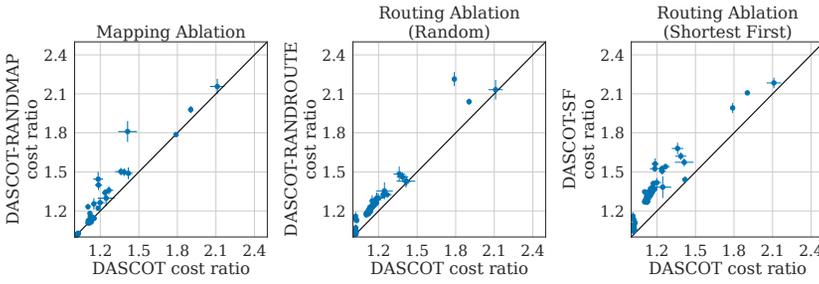


Fig. 17. Effects of mapping and routing

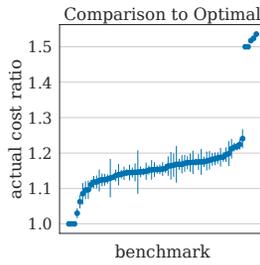


Fig. 18. Optimality on the Compact architecture

applying dependency-aware routing for a mean relative improvement of 6%. Similar findings hold when we compare DASCOT to choosing a routing order by sorting gates according to routing path length and criticality. We show the results for the shortest-first routing order, which is the approach of EDPC [8], in Fig. 17(right). The plot for the critical-first sorting order closely resembles this one and is included in the extended version of this paper [46].

RQ2 Summary: Optimization at both stages contributes to the performance of our algorithm. However, more consistent gains can be found at routing phase, where optimizing routing order improves results on 90% of circuits, and by 6% on average.

(RQ3) Comparison to Optimal

We also assess how well DASCOT approximates the optimal solution using the baseline described in Section 5. Our implementation generates the constraints, then queries the SAT solver CaDiCal [10] via the PySAT toolkit [32]. We focus on the Compact architecture, as DASCOT generally reaches the depth lower-bound on the Square Sparse architecture, so we know the solutions are optimal.

In Fig. 18, we compare DASCOT to the optimal solution in 69 cases where we are able to obtain one within the time and memory bounds. Circuits are sorted by their *actual cost ratio*, which replaces the theoretical, architecture-independent lower bound from the cost ratio with the true optimal solution for the Compact architecture. For 64 of the 69 circuits, the mean actual cost ratio is less than 1.25, meaning the mean solution produced by DASCOT is within 25% of optimal. A common characteristic shared by the five outlier circuits is relatively high density. Such circuits remain challenging for DASCOT despite improvements over the state of the art.

RQ3 Summary: For small circuits where optimal solving is feasible, we find that DASCOT finds solutions that are within 1.25x optimal for all but a few examples.

(RQ4) Optimality-Scalability Tradeoff

We now study how the runtime of our algorithm scales with the complexity of the SCMR problem. In particular, we are interested in the routing algorithm because mapping can be terminated at any point and our ablation study suggests most of the improvements in solution quality are found at the routing stage. The complexity of an SCMR problem is multidimensional, depending on parameters which include the number of qubits in the input circuit, the depth of the circuit, and the density of the topological layers. We highlight two circuit classes with contrasting characteristics: Bernstein-Vazirani circuits and Quantum Fourier Transform circuits.

Bernstein-Vazirani circuits have low depth and each layer is minimal, consisting of a single CNOT gate. Because of this structure, our approach can route Bernstein-Vazirani circuits with thousands of qubits, as shown in Fig. 19(left). On the other hand, QFT circuits of the same qubit count have many more layers, with several gates per layer. Thus, the problem complexity grows much more rapidly with respect to qubit count (see Fig. 19(right)), such that the largest circuit routed within an hour has 170 qubits on the Compact architecture and 200 qubits on the Square Sparse architecture.

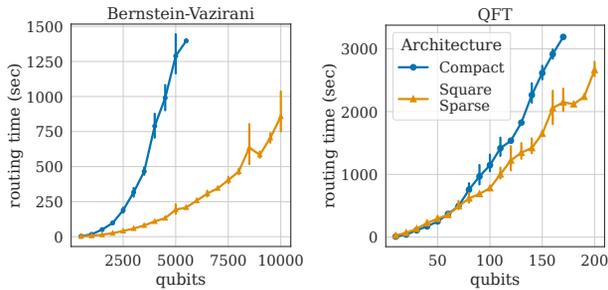


Fig. 19. Runtime scaling of DASCOT

However, our approach is flexible and can terminate more quickly by searching for fewer iterations. In Fig. 20, we compare the solution quality and runtime of three routing algorithms: DASCOT and DASCOT-RANDROUTE from our ablation study (the latter is abbreviated to RAND in the legend) and an intermediate configuration, LIMITED, which performs 1/2 as many search iterations as DASCOT. Benchmarks are sorted along the x -axis in ascending order with respect to LIMITED. We see that LIMITED finds solutions with cost ratios near the midpoint between DASCOT-RANDROUTE (2% better than DASCOT-RANDROUTE and 3% worse than DASCOT on average) while terminating 41% faster than DASCOT on average.

RQ4 Summary: DASCOT can be applied to Bernstein-Vazirani circuits with up to 10,000 qubits and QFT circuits up to 200. Additionally, applying a limited version of DASCOT decreases runtime by 41%, incurring a only a 3% cost in terms of solution quality.

7 Discussion and Related Work

Other allocation and disjoint paths problems. Surface code mapping and routing is an assignment of program variables to limited physical resources such that the runtime efficiency is maximized, much like the classical problem of register allocation [14, 17, 48]. Disjoint paths problems like the constraints on a valid gate route have also been studied in other contexts including

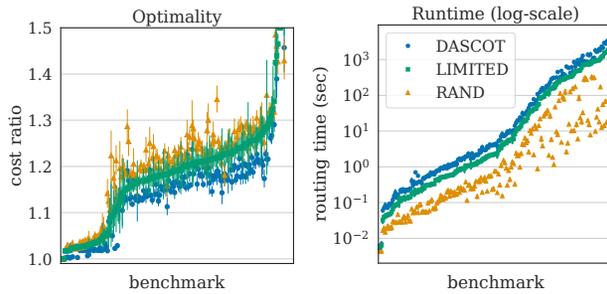


Fig. 20. Trading compute for optimality

VLSI routing [2, 38] and all-optical networks [1, 5, 50]. Inspired by these applications, the study of approximation algorithms for disjoint path problems is an active area with both hardness results [3, 21] and approximation algorithms on planar and grid graphs [15, 16]. However, a distinguishing feature of this particular disjoint path problem is the dependency between paths, forcing the routing of certain pairs in a particular order.

NISQ qubit mapping and routing. A wide array of work has targeted the qubit mapping and CNOT routing problem for noisy intermediate-scale quantum (NISQ) computers without error-correction. In the NISQ setting, CNOT gates are only executable between qubits which are mapped to adjacent vertices on the architecture graph. NISQ mapping and routing tools use SWAP gates to update the qubit map, moving qubits to adjacent locations. The constraints of SCMR are fundamentally different: we can perform two qubit gates between any pair of qubits in constant time as long as there is an appropriate routing path available. Nevertheless, we can leverage certain insights developed in the NISQ setting. For example, our layered interaction graph approach for mapping is an extension of NISQ algorithms those that construct qubit maps based on interaction graphs [19, 57]. Also, our SAT encoding for the optimal baseline shares structure with encodings of the NISQ problem in SAT, SMT, or MAXSAT [42, 45, 53, 60, 63, 65], though the constraints describing valid paths are unique to SCMR and thus to our encoding.

Defect-braid routing. Some recent work in compilation for surface code quantum devices has addressed similar problems to the surface code mapping and routing problem. Javadi-Abhari et al. [33] identifies CNOT contention as an important architectural design factor and develops heuristics for mapping qubits and scheduling CNOT gates implemented as defect braids. Autobraid [31] builds on this work with a new algorithm for mapping that minimizes the number of large groups of CNOT gates with overlapping bounding boxes and a stack-based routing algorithm. However, neither considers τ gates, multiple architectures, or the constraints of lattice surgery CNOT gates.

Autobraid includes a procedure for inserting SWAP operations to shuffle the qubit map when a small proportion of front layer of gates can be routed in the same time step. In principle, such shuffling can be incorporated into DASCOT. However, SWAP gates are expensive, decomposing into three CNOT gates of alternating orientation (the control of one CNOT is the target of the next). Because CNOT gates of opposite orientation must be routed along different paths, SWAP gates occupy many vertices and are in parallel with other gates. Thus, transitioning to a new mapping often has a net negative effect as the added steps for SWAP gates outweighs the benefit of the new mapping.

Lattice surgery compilation. For the lattice surgery setting, LSC [62] provides a scalable framework for translation of quantum gates to lattice surgery operations, but does not optimize the mapping or routing solution. EDPC [8] considers a shortest-first routing algorithm, but assumes

a given fixed map that places qubits in a “sparse” structure. None of the above work provides an optimal solver for the compilation problem they address. Lao et al. [40] does solve a mapping and routing problem optimally with integer linear programming, but this formulation does not permit long-range CNOT gates. LaSynth [61] uses a SAT solver to optimize lattice surgery representations of small subroutines (5-20 qubits and 10-100 operations) at a lower level of abstraction.

An alternative proposed compilation pipeline serializes a circuit to a sequence of Pauli product rotations [9, 43]. In this setting, TopQAD [55] applies an algorithm similar to LSC: performing no mapping optimization and routing gates in a random order. The only difference between this greedy routing algorithm and DASCOT-RANDROUTE is in the specifics of the routed operations. Pauli product rotations are implemented via routing trees with an interacting qubit at each leaf, rather than routing paths with an interacting qubit at each endpoint. However, sequential Pauli compilation can lead to prohibitively high run times, converting a circuit with k parallel operations to a compiled circuit of depth k [8].

8 Conclusions

In this paper, we have tackled the surface code mapping and routing problem, a critical problem in compilation for emerging practical quantum computers. We developed an algorithm to deliver high-quality solutions in reasonable time via simulated annealing and dependency-awareness. In future work, we plan to extend our model and algorithms to capture a broader class of quantum architectures. For example, we could support recent designs for heterogeneous multi-chip fault-tolerant quantum architectures [58, 59]. The mapping and routing problem in this case is nuanced because not all CNOT gates are equivalent. A cross-chip CNOT gate is possible, but more costly.

Data-Availability Statement

The software and benchmark suite that supports our evaluation in Section 6 is available in a public archive [47]. This software artifact includes a Docker image packaging the source code for DASCOT and the baseline algorithms, along with scripts for reproducing our empirical results and generating plots matching those included in the paper.

Acknowledgments

We thank the anonymous reviewers for their insightful feedback and suggestions. We are grateful to the CHTC team for maintaining the computing resources which enabled our empirical evaluation [13]. This work is supported by NSF grants #1652140 and #2212232 and awards from Meta and Amazon. This research is also partially supported by the OVCRGE at the University of Wisconsin-Madison with funding from the Wisconsin Alumni Research Foundation.

References

- [1] Alok Aggarwal, Amotz Bar-Noy, Don Coppersmith, Rajiv Ramaswami, Baruch Schieber, and Madhu Sudan. 1994. Efficient Routing and Scheduling Algorithms for Optical Networks. In *Proceedings of the Fifth Annual ACM-SIAM Symposium on Discrete Algorithms* (Arlington, Virginia, USA) (SODA '94). Society for Industrial and Applied Mathematics, USA, 412–423.
- [2] Alok Aggarwal, Jon Kleinberg, and David P. Williamson. 2000. Node-Disjoint Paths on the Mesh and a New Trade-Off in VLSI Layout. *SIAM J. Comput.* 29, 4 (2000), 1321–1333. <https://doi.org/10.1137/S0097539796312733>
- [3] M. Andrews, J. Chuzhoy, Sanjeev Khanna, and L. Zhang. 2005. Hardness of the undirected edge-disjoint paths problem with congestion. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*. 226–241. <https://doi.org/10.1109/SFCS.2005.41>
- [4] Carlos Ansótegui and Felip Manyà. 2004. Mapping problems with finite-domain variables into problems with boolean variables. 1–15. https://doi.org/10.1007/11527695_1

- [5] Yonatan Aumann and Yuval Rabani. 1995. Improved Bounds for All Optical Routing. In *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms* (San Francisco, California, USA) (SODA '95). Society for Industrial and Applied Mathematics, USA, 567–576.
- [6] Olivier Bailleux and Yacine Boufkhad. 2003. Efficient CNF encoding of Boolean cardinality constraints. In *Proceedings of the 9th International Conference on Principles and Practice of Constraint Programming* (Kinsale, Ireland) (CP'03). Springer-Verlag, Berlin, Heidelberg, 108–122. https://doi.org/10.1007/978-3-540-45193-8_8
- [7] Ethan Bernstein and Umesh Vazirani. 1997. Quantum Complexity Theory. *SIAM J. Comput.* 26, 5 (1997), 1411–1473. <https://doi.org/10.1137/S0097539796300921> arXiv:<https://doi.org/10.1137/S0097539796300921>
- [8] Michael Beverland, Vadym Kliuchnikov, and Eddie Schoute. 2022. Surface Code Compilation via Edge-Disjoint Paths. *PRX Quantum* 3 (May 2022), 020342. Issue 2. <https://doi.org/10.1103/PRXQuantum.3.020342>
- [9] Michael E. Beverland, Prakash Murali, Matthias Troyer, Krysta M. Svore, Torsten Hoefler, Vadym Kliuchnikov, Guang Hao Low, Mathias Soeken, Aarthi Sundaram, and Alexander Vaschillo. 2022. Assessing requirements to scale to practical quantum advantage. (11 2022). <https://doi.org/10.48550/arXiv.2211.07629> arXiv:2211.07629 [quant-ph]
- [10] Armin Biere, Katalin Fazekas, Mathias Fleury, and Maximillian Heisinger. 2020. CaDiCaL, Kissat, Paracooba, Plingeling and Treengeling Entering the SAT Competition 2020. In *Proc. of SAT Competition 2020 – Solver and Benchmark Descriptions (Department of Computer Science Report Series B, Vol. B-2020-1)*, Tomas Balyo, Nils Froyleyks, Marijn Heule, Markus Iser, Matti Järvisalo, and Martin Suda (Eds.). University of Helsinki, 51–53.
- [11] Dolev Bluvstein, Simon J. Evered, Alexandra A. Geim, Sophie H. Li, Hengyun Zhou, Tom Manovitz, Sepehr Ebadi, Madelyn Cain, Marcin Kalinowski, Dominik Hangleiter, J. Pablo Bonilla Ataides, Nishad Maskara, Iris Cong, Xun Gao, Pedro Sales Rodriguez, Thomas Karolyshyn, Giulia Semeghini, Michael J. Gullans, Markus Greiner, Vladan Vuletić, and Mikhail D. Lukin. 2023. Logical quantum processor based on reconfigurable atom arrays. *Nature* 626, 7997 (Dec. 2023), 58–65. <https://doi.org/10.1038/s41586-023-06927-3>
- [12] Sergey Bravyi and Alexei Kitaev. 2005. Universal quantum computation with ideal Clifford gates and noisy ancillas. *Physical Review A* 71, 2 (Feb. 2005). <https://doi.org/10.1103/physreva.71.022316>
- [13] Center for High Throughput Computing. 2006. Center for High Throughput Computing. <https://doi.org/10.21231/GNT1-HW21>
- [14] G. J. Chaitin. 1982. Register Allocation & Spilling via Graph Coloring. *SIGPLAN Not.* 17, 6 (jun 1982), 98–101. <https://doi.org/10.1145/872726.806984>
- [15] Julia Chuzhoy and David H. K. Kim. 2015. On Approximating Node-Disjoint Paths in Grids. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2015) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 40)*, Naveen Garg, Klaus Jansen, Anup Rao, and José D. P. Rolim (Eds.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 187–211. <https://doi.org/10.4230/LIPIcs.APPROX-RANDOM.2015.187>
- [16] Julia Chuzhoy, David H. K. Kim, and Shi Li. 2016. Improved Approximation for Node-Disjoint Paths in Planar Graphs. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing* (Cambridge, MA, USA) (STOC '16). Association for Computing Machinery, New York, NY, USA, 556–569. <https://doi.org/10.1145/2897518.2897538>
- [17] Keith D. Cooper and Linda Torczon. 2023. Chapter 13 - Register Allocation. In *Engineering a Compiler (Third Edition)* (third edition ed.), Keith D. Cooper and Linda Torczon (Eds.). Morgan Kaufmann, Philadelphia, 663–712. <https://doi.org/10.1016/B978-0-12-815412-0.00019-X>
- [18] Don Coppersmith. 2002. An approximate Fourier transform useful in quantum factoring. *arXiv preprint quant-ph/0201067* (2002). <https://doi.org/10.48550/arXiv.quant-ph/0201067>
- [19] Alexander Cowtan, Silas Dilkes, Ross Duncan, Alexandre Krajenbrink, Will Simmons, and Seyon Sivarajah. 2019. On the Qubit Routing Problem. *Leibniz Int. Proc. Inf.* 135 (2019), 5:1–5:32. <https://doi.org/10.4230/LIPIcs.TQC.2019.5> arXiv:1902.08091 [quant-ph]
- [20] M. P. da Silva, C. Ryan-Anderson, J. M. Bello-Rivas, A. Chernoguzov, J. M. Dreiling, C. Foltz, F. Frachon, J. P. Gaebler, T. M. Gatterman, L. Grans-Samuelsson, D. Hayes, N. Hewitt, J. Johansen, D. Lucchetti, M. Mills, S. A. Moses, B. Neyenhuis, A. Paz, J. Pino, P. Siegfried, J. Strabley, A. Sundaram, D. Tom, S. J. Wernli, M. Zanner, R. P. Stutz, and K. M. Svore. 2024. Demonstration of logical qubits and repeated error correction with better-than-physical error rates. <https://doi.org/10.48550/arXiv.2404.02280> arXiv:2404.02280
- [21] Thomas Erlebach. 2006. Approximation Algorithms for Edge-Disjoint Paths and Unsplittable Flow. In *Efficient Approximation and Online Algorithms - Recent Progress on Classical Combinatorial Optimization Problems and New Applications*, Evripidis Bampis, Klaus Jansen, and Claire Kenyon (Eds.). Lecture Notes in Computer Science, Vol. 3484. Springer, 97–134. https://doi.org/10.1007/11671541_4
- [22] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. 2014. A Quantum Approximate Optimization Algorithm. <https://doi.org/10.48550/arXiv.1411.4028> arXiv:1411.4028

- [23] Austin G. Fowler and Craig Gidney. 2018. Low overhead quantum computation using lattice surgery. *arXiv: Quantum Physics* (2018). <https://doi.org/10.48550/arXiv.1808.06709>
- [24] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. 2012. Surface codes: Towards practical large-scale quantum computation. *Physical Review A* 86, 3 (sep 2012). <https://doi.org/10.1103/physreva.86.032324>
- [25] Ian P. Gent and Peter Nightingale. 2004. A new encoding of All Different into SAT. In *International Workshop on Modelling and Reformulating Constraint Satisfaction Problems*, Vol. 3. 95–110.
- [26] Google Quantum AI. 2023. Suppressing quantum errors by scaling a surface code logical qubit. *Nature* 614, 7949 (Feb. 2023), 676–681. <https://doi.org/10.1038/s41586-022-05434-1>
- [27] Google Quantum AI and Collaborators. 2024. Quantum error correction below the surface code threshold. *Nature* (Dec. 2024). <https://doi.org/10.1038/s41586-024-08449-y>
- [28] Alexander S. Green, Peter LeFanu Lumsdaine, Neil J. Ross, Peter Selinger, and Benoit Valiron. 2013. Quipper: a scalable quantum programming language. *SIGPLAN Not.* 48, 6 (June 2013), 333–342. <https://doi.org/10.1145/2499370.2462177>
- [29] Lov K. Grover. 1996. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing* (Philadelphia, Pennsylvania, USA) (STOC '96). Association for Computing Machinery, New York, NY, USA, 212–219. <https://doi.org/10.1145/237814.237866>
- [30] Clare Horsman, Austin G Fowler, Simon Devitt, and Rodney Van Meter. 2012. Surface code quantum computing by lattice surgery. *New Journal of Physics* 14, 12 (dec 2012), 123011. <https://doi.org/10.1088/1367-2630/14/12/123011>
- [31] Fei Hua, Yanhao Chen, Yuwei Jin, Chi Zhang, Ari Hayes, Youtao Zhang, and Eddy Z. Zhang. 2021. AutoBraid: A Framework for Enabling Efficient Surface Code Communication in Quantum Computing. In *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture* (Virtual Event, Greece) (MICRO '21). Association for Computing Machinery, New York, NY, USA, 925–936. <https://doi.org/10.1145/3466752.3480072>
- [32] Alexey Ignatiev, Antonio Morgado, and Joao Marques-Silva. 2018. PySAT: A Python Toolkit for Prototyping with SAT Oracles. In *SAT*. 428–437. https://doi.org/10.1007/978-3-319-94144-8_26
- [33] Ali Javadi-Abhari, Pranav Gokhale, Adam Holmes, Diana Franklin, Kenneth R. Brown, Margaret Martonosi, and Frederic T. Chong. 2017. Optimized Surface Code Communication in Superconducting Quantum Computers. In *Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture* (Cambridge, Massachusetts) (MICRO-50 '17). Association for Computing Machinery, New York, NY, USA, 692–705. <https://doi.org/10.1145/3123939.3123949>
- [34] Ali Javadi-Abhari, Shruti Patil, Daniel Kudrow, Jeff Heckey, Alexey Lvov, Frederic T. Chong, and Margaret Martonosi. 2014. Scaffold: A Framework for Compilation and Analysis of Quantum Computing Programs. In *Proceedings of the 11th ACM Conference on Computing Frontiers* (Cagliari, Italy) (CF '14). Association for Computing Machinery, New York, NY, USA, Article 1, 10 pages. <https://doi.org/10.1145/2597917.2597939>
- [35] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. 1983. Optimization by Simulated Annealing. *Science* 220, 4598 (1983), 671–680. <https://doi.org/10.1126/science.220.4598.671>
- [36] Emanuel Knill, Raymond Laflamme, and Wojciech H. Zurek. 1998. Resilient quantum computation: error models and thresholds. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* 454, 1969 (jan 1998), 365–384. <https://doi.org/10.1098/rspa.1998.0166>
- [37] Stavros G. Kolliopoulos and Clifford Stein. 1997. Approximating Disjoint-Path Problems Using Greedy Algorithms and Packing Integer Programs. In *Conference on Integer Programming and Combinatorial Optimization*. <https://doi.org/10.1007/s10107-002-0370-6>
- [38] M.R Kramer and J van Leeuwen. 1984. The complexity of wire-routing and finding minimum area layouts for arbitrary vlsi circuits. *Advances in Computing Research* 2 (1984), 020342.
- [39] J M Kreikebaum, K P O'Brien, A Morvan, and I Siddiqi. 2020. Improving wafer-scale Josephson junction resistance variation in superconducting quantum coherent circuits. *Superconductor Science and Technology* 33, 6 (apr 2020), 06LT02. <https://doi.org/10.1088/1361-6668/ab8617>
- [40] L Lao, B van Wee, I Ashraf, J van Someren, N Khammassi, K Bertels, and C G Almudever. 2018. Mapping of lattice surgery-based quantum circuits on surface code architectures. *Quantum Science and Technology* 4, 1 (sep 2018), 015005. <https://doi.org/10.1088/2058-9565/aadd1a>
- [41] Sophia Fuhui Lin, Joshua Vizslai, Kaitlin N. Smith, Gokul Subramanian Ravi, Charles Yuan, Frederic T. Chong, and Benjamin J. Brown. 2024. Codesign of quantum error-correcting codes and modular chiplets in the presence of defects. In *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2* (La Jolla, CA, USA) (ASPLOS '24). Association for Computing Machinery, New York, NY, USA, 216–231. <https://doi.org/10.1145/3620665.3640362>
- [42] Wan-Hsuan Lin, Jason Kimko, Bochen Tan, Nikolaj Björner, and Jason Cong. 2023. Scalable Optimal Layout Synthesis for NISQ Quantum Processors. In *2023 60th ACM/IEEE Design Automation Conference (DAC)*. 1–6. <https://doi.org/10.1109/DAC56929.2023.10247760>

- [43] Daniel Litinski. 2019. A Game of Surface Codes: Large-Scale Quantum Computing with Lattice Surgery. *Quantum* 3 (March 2019), 128. <https://doi.org/10.22331/q-2019-03-05-128>
- [44] Daniel Litinski and Felix von Oppen. 2018. Lattice Surgery with a Twist: Simplifying Clifford Gates of Surface Codes. *Quantum* 2 (May 2018), 62. <https://doi.org/10.22331/q-2018-05-04-62>
- [45] Abtin Molavi, Amanda Xu, Martin Diges, Lauren Pick, Swamit S. Tannu, and Aws Albarghouthi. 2022. Qubit Mapping and Routing via MaxSAT. *2022 55th IEEE/ACM International Symposium on Microarchitecture (MICRO)* (2022), 1078–1091. <https://doi.org/10.1109/MICRO56248.2022.00077>
- [46] Abtin Molavi, Amanda Xu, Swamit Tannu, and Aws Albarghouthi. 2024. Dependency-Aware Compilation for Surface Code Quantum Architectures. <https://doi.org/10.48550/arXiv.2311.18042> arXiv:2311.18042 [quant-ph]
- [47] Abtin Molavi, Amanda Xu, Swamit Tannu, and Aws Albarghouthi. 2025. Artifact for "Dependency-Aware Compilation for Surface Code Quantum Architectures" (DASCOT). <https://doi.org/10.5281/zenodo.14934044>
- [48] Massimiliano Poletto and Vivek Sarkar. 1999. Linear Scan Register Allocation. *ACM Trans. Program. Lang. Syst.* 21, 5 (sep 1999), 895–913. <https://doi.org/10.1145/330249.330250>
- [49] Steven Prestwich. 2009. CNF encodings. *Frontiers in Artificial Intelligence and Applications* 185 (01 2009). <https://doi.org/10.3233/978-1-58603-929-5-75>
- [50] Prabhakar Raghavan and Eli Upfal. 1994. Efficient Routing in All-Optical Networks. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing* (Montreal, Quebec, Canada) (*STOC '94*). Association for Computing Machinery, New York, NY, USA, 134–143. <https://doi.org/10.1145/195058.195119>
- [51] Ben W. Reichardt, David Aasen, Rui Chao, Alex Chernoguzov, Wim van Dam, John P. Gaebler, Dan Gresh, Dominic Lucchetti, Michael Mills, Steven A. Moses, Brian Neyenhuis, Adam Paetznick, Andres Paz, Peter E. Siegfried, Marcus P. da Silva, Krysta M. Svore, Zhenghan Wang, and Matt Zanner. 2024. Demonstration of quantum computation and error correction with a tesseract code. <https://doi.org/10.48550/arXiv.2409.04628> arXiv:2409.04628 [quant-ph]
- [52] C. Ryan-Anderson, J. G. Bohnet, K. Lee, D. Gresh, A. Hankin, J. P. Gaebler, D. Francois, A. Chernoguzov, D. Lucchetti, N. C. Brown, T. M. Gatterman, S. K. Halit, K. Gilmore, J. A. Gerber, B. Neyenhuis, D. Hayes, and R. P. Stutz. 2021. Realization of Real-Time Fault-Tolerant Quantum Error Correction. *Phys. Rev. X* 11 (Dec 2021), 041058. Issue 4. <https://doi.org/10.1103/PhysRevX.11.041058>
- [53] Irfansha Shaik and Jaco van de Pol. 2024. Optimal Layout Synthesis for Deep Quantum Circuits on NISQ Processors with 100+ Qubits. In *27th International Conference on Theory and Applications of Satisfiability Testing (SAT 2024) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 305)*, Supratik Chakraborty and Jie-Hong Roland Jiang (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 26:1–26:18. <https://doi.org/10.4230/LIPIcs.SAT.2024.26>
- [54] P.W. Shor. 1994. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*. 124–134. <https://doi.org/10.1109/SFCS.1994.365700>
- [55] Allyson Silva, Xiangyi Zhang, Zak Webb, Mia Kramer, Chan-Woo Yang, Xiao Liu, Jessica Lemieux, Ka-Wai Chen, Artur Scherer, and Pooya Ronagh. 2024. Multi-qubit Lattice Surgery Scheduling. *Schloss Dagstuhl – Leibniz-Zentrum für Informatik*. <https://doi.org/10.4230/LIPIcs.TQC.2024.1>
- [56] Carsten Sinz. 2005. Towards an Optimal CNF Encoding of Boolean Cardinality Constraints. In *Principles and Practice of Constraint Programming - CP 2005*, Peter van Beek (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 827–831. https://doi.org/10.1007/11564751_73
- [57] Marcos Yukio Siraichi, Vinicius Fernandes dos Santos, Caroline Collange, and Fernando Magno Quintão Pereira. 2019. Qubit allocation as a combination of subgraph isomorphism and token swapping. *Proc. ACM Program. Lang.* 3, OOPSLA, Article 120 (Oct. 2019), 29 pages. <https://doi.org/10.1145/3360546>
- [58] K. N. Smith, G. Ravi, J. M. Baker, and F. T. Chong. 2022. Scaling Superconducting Quantum Computers with Chiplet Architectures. In *2022 55th IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE Computer Society, Los Alamitos, CA, USA, 1092–1109. <https://doi.org/10.1109/MICRO56248.2022.00078>
- [59] Samuel Stein, Sara Sussman, Teague Tomesh, Charles Guinn, Esin Tureci, Sophia Fuhui Lin, Wei Tang, James Ang, Srivatsan Chakram, Ang Li, Margaret Martonosi, Fred Chong, Andrew A. Houck, Isaac L. Chuang, and Michael Demarco. 2023. HetArch: Heterogeneous Microarchitectures for Superconducting Quantum Systems. In *Proceedings of the 56th Annual IEEE/ACM International Symposium on Microarchitecture* (Toronto, ON, Canada) (*MICRO '23*). Association for Computing Machinery, New York, NY, USA, 539–554. <https://doi.org/10.1145/3613424.3614300>
- [60] Bochen Tan and Jason Cong. 2020. Optimal Layout Synthesis for Quantum Computing. In *Proceedings of the 39th International Conference on Computer-Aided Design* (Virtual Event, USA) (*ICCAD '20*). Association for Computing Machinery, New York, NY, USA, Article 137, 9 pages. <https://doi.org/10.1145/3400302.3415620>
- [61] Daniel Bochen Tan, Murphy Yuezhen Niu, and Craig Gidney. 2024. A SAT Scalpel for Lattice Surgery: Representation and Synthesis of Subroutines for Surface-Code Fault-Tolerant Quantum Computing. In *2024 ACM/IEEE 51st Annual International Symposium on Computer Architecture (ISCA)*. IEEE Computer Society, Los Alamitos, CA, USA, 325–339. <https://doi.org/10.1109/ISCA59077.2024.00032>

- [62] George Watkins, Hoang Minh Nguyen, Keelan Watkins, Steven Pearce, Hoi-Kwan Lau, and Alexandru Paler. 2024. A High Performance Compiler for Very Large Scale Surface Code Computations. *Quantum* 8, 1354. <https://doi.org/10.22331/q-2024-05-22-1354> arXiv:2302.02459 [quant-ph]
- [63] Robert Wille, Lukas Burgholzer, and Alwin Zulehner. 2019. Mapping Quantum Circuits to IBM QX Architectures Using the Minimal Number of SWAP and H Operations. In *Proceedings of the 56th Annual Design Automation Conference 2019 (Las Vegas, NV, USA) (DAC '19)*. Association for Computing Machinery, New York, NY, USA, Article 142, 6 pages. <https://doi.org/10.1145/3316781.3317859>
- [64] Robert Wille, Daniel Große, Lisa Teuber, Gerhard W. Dueck, and Rolf Drechsler. 2008. RevLib: An Online Resource for Reversible Functions and Reversible Circuits. In *38th International Symposium on Multiple Valued Logic (ismvl 2008)*. 220–225. <https://doi.org/10.1109/ISMVL.2008.43>
- [65] Jiong Yang, Yaroslav A. Kharkov, Yunong Shi, Marijn J. H. Heule, and Bruno Dutertre. 2024. Quantum Circuit Mapping Based on Incremental and Parallel SAT Solving. In *27th International Conference on Theory and Applications of Satisfiability Testing (SAT 2024) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 305)*, Supratik Chakraborty and Jie-Hong Roland Jiang (Eds.). Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 29:1–29:18. <https://doi.org/10.4230/LIPIcs.SAT.2024.29>
- [66] Youwei Zhao, Yangsen Ye, He-Liang Huang, Yiming Zhang, Dachao Wu, Huijie Guan, Qingling Zhu, Zuolin Wei, Tan He, Sirui Cao, Fusheng Chen, Tung-Hsun Chung, Hui Deng, Daojin Fan, Ming Gong, Cheng Guo, Shaojun Guo, Lianchen Han, Na Li, Shaowei Li, Yuan Li, Futian Liang, Jin Lin, Haoran Qian, Hao Rong, Hong Su, Lihua Sun, Shiyu Wang, Yulin Wu, Yu Xu, Chong Ying, Jiale Yu, Chen Zha, Kaili Zhang, Yong-Heng Huo, Chao-Yang Lu, Cheng-Zhi Peng, Xiaobo Zhu, and Jian-Wei Pan. 2022. Realization of an Error-Correcting Surface Code with Superconducting Qubits. *Phys. Rev. Lett.* 129 (Jul 2022), 030501. Issue 3. <https://doi.org/10.1103/PhysRevLett.129.030501>
- [67] Mingzheng Zhu, Hao Fu, Jun Wu, Chi Zhang, Wei Xie, and Xiang-Yang Li. 2024. Ecmass: Efficient Circuit Mapping and Scheduling for Surface Code. In *2024 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*. IEEE Computer Society, Los Alamitos, CA, USA, 158–169. <https://doi.org/10.1109/CGO57630.2024.10444874>
- [68] Alwin Zulehner, Alexandru Paler, and Robert Wille. 2018. An efficient methodology for mapping quantum circuits to the IBM QX architectures. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 38, 7, 1226–1236. <https://doi.org/10.1109/TCAD.2018.2846658>

Received 2024-10-11; accepted 2025-02-18