Optimal Layout Synthesis for Quantum Computing

Bochen Tan

University of California, Los Angeles
bctan@cs.ucla.edu

Jason Cong

University of California, Los Angeles
cong@cs.ucla.edu

Abstract

Recent years have witnessed the fast development of quantum computing. Researchers around the world are eager to run larger and larger quantum algorithms that promise speedups impossible to any classical algorithm. However, the available quantum computers are still volatile and error-prone. Thus, layout synthesis, which transforms quantum programs to meet these hardware limitations, is a crucial step in the realization of quantum computing. In this paper, we present two synthesizers, one optimal and one approximate but nearly optimal. Although a few optimal approaches to this problem have been published, our optimal synthesizer explores a larger solution space, thus is optimal in a stronger sense. In addition, it reduces time and space complexity exponentially compared to some leading optimal approaches. The key to this success is a more efficient spacetime-based variable encoding of the layout synthesis problem as a mathematical programming problem. By slightly changing our formulation, we arrive at an approximate synthesizer that is even more efficient and outperforms some leading heuristic approaches, in terms of additional gate cost, by up to 100%, and also fidelity by up to 10x on a comprehensive set of benchmark programs and architectures. For a specific family of quantum programs named QAOA, which is deemed to be a promising application for near-term quantum computers, we further adjust the approximate synthesizer by taking commutation into consideration, achieving up to 75% reduction in depth and up to 65% reduction in additional cost compared to the tool used in a leading QAOA study.

CCS Concepts: • Computer systems organization → Quantum computing; • Hardware → Physical synthesis.

Keywords: quantum computing, layout synthesis, allocation, placement, scheduling, mapping

ACM Reference Format:

Bochen Tan and Jason Cong. 2020. Optimal Layout Synthesis for Quantum Computing. In *IEEE/ACM International Conference on Computer-Aided Design (ICCAD '20)*, November 2–5, 2020, Virtual Event, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3400302.3415620>

1 Introduction

For years, quantum algorithms have been shown theoretically to hold significant advantages over classical algorithms on some algebraic problems like factoring large numbers [20], quantum chemistry [5], machine learning [7], and so on. Recently, the first experimental proof of quantum computational advantage on certain sampling problems was published [3].

To implement a quantum algorithm on an actual quantum computer, *logic synthesis* and *layout synthesis* for quantum computing (QC) are inevitable. A classical algorithm is a series of operations on a set of bits; likewise, a quantum algorithm is a series of operations, i.e., *quantum gates*, on a set of quantum bits, i.e., *qubits*. A realistic *quantum computer architecture* only supports gates in its *gate library*. To implement a quantum algorithm, the original quantum gates have to be translated into gates within this library. This process is called logic synthesis for quantum computing. It has been shown that all quantum algorithms can be decomposed to quantum gates acting on a single qubit or two qubits [17], so usually the output of logic synthesis is a list of single-qubit gates and two-qubit gates.

After logic synthesis, the qubits and quantum gates are still purely logical. Next, a mapping from such *logical qubits* to the *physical qubits* on QC architecture and a schedule which decides when every quantum gate is executed have to be provided. (Note that our notion of logical qubit is different from that in the quantum error correction.) The task of deriving such information is called layout synthesis for quantum computing (LSQC) [23]. Other names have also been used to refer to this task, e.g., quantum circuit compilation [1, 8], transpilation (in Qiskit¹), placement [14, 19], mapping [6, 13, 15, 25, 28], qubit allocation [21], and routing [22]. We choose to use the term ‘layout synthesis’ as it involves both placement and scheduling. There is an important issue to resolve in LSQC. Every two-qubit gate requires a connection between its two qubits. An algorithm designer may assume the logical qubits to have all-to-all connectivity, but this is not the case for the physical qubits in many QC architectures, e.g., the superconducting QC systems at Google [3]. As a result, every QC architecture comes with a *coupling graph* specifying which physical qubit pairs are connected. If the pair of logical qubits of a two-qubit gate is mapped to a disconnected physical qubit pair on the coupling graph, this gate will not be executable. Luckily, there is a special type of gate, SWAP, that can exchange the mapping of two logical qubits. Thus, with a series of SWAP gates, two logical qubits can reach two connected physical qubits, but this comes at some additional cost. The decisions of when and where to apply these SWAP gates are made in LSQC.

A recent study [23] revealed large optimality gaps of several leading LSQC tools with a set of benchmarks named QUEKO that have known optimal solutions. Therefore, despite the NP-completeness of the problem [14, 21, 23], optimal approaches to LSQC are still worthy of research because they provide the baseline of measuring the available LSQC solutions. There have been a few previous works on exact or optimal LSQC [6, 19, 21, 25, 27]. However, they either just focus on specific coupling graphs, or has some implicit and unnecessary constraints. As for the heuristic approaches [8, 13–15, 21, 22, 24, 28], the optimality gaps suggest that they still have a lot of room for improvement.

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

ICCAD '20, November 2–5, 2020, Virtual Event, USA

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8026-3/20/11.

<https://doi.org/10.1145/3400302.3415620>

¹<https://qiskit.org/>

In this paper, we present two layout synthesizers and compare them with leading related works on a comprehensive set of benchmark quantum programs and architectures. The optimal layout synthesizer for quantum computing (OLSQ) formulates LSQC as a satisfiability modulo theories optimization problem (SMT). This problem is then passed to Z3 SMT solver [9] which guarantees optimal solutions. Compared to Wille et al. [25], a leading optimal approach which also formulates LSQC to SMT, we provide exponential reduction on the number of variables, resulting in orders-of-magnitude reductions in runtime and the memory usage. The key to this improvement is to represent the solution space with mapping variables, not with mapping transformation variables as in [21, 25]. Also, by relaxing the gate placement under the *dependency* constraints, OLSQ explores a larger solution space than other gate-by-gate [21, 25] and level-by-level [6] “optimal” approaches, so its results are sometimes even better. Furthermore, by eliminating redundant mapping variables between *transitions*, i.e., when the mapping changes due to some SWAP gates, we significantly reduce the number of variables and derive a highly scalable approximate synthesizer, transition-based-OLSQ (TB-OLSQ). Evaluation results show that TB-OLSQ is nearly optimal and is able to increase fidelity by 1.30x in geomean, compared to TriQ [15], a leading academic work, and reduce cost by 69.2% in geomean, compared to t|ket> [22], a leading industry work. Finally, given that quantum approximate optimization (QAOA) [4, 10, 12, 18] is very much in the spotlight due to its feasibility for near-term quantum computers, we study the LSQC of QAOA. The structure of gates in QAOA makes it an excellent showcase for the power of our transition-based model. Moreover, with additional prior knowledge on commutation relations in QAOA, we improve TB-OLSQ for QAOA, denoted as QAOA-OLSQ. Compared to t|ket>, which was utilized in a leading QAOA experiment study [4], QAOA-OLSQ is able to reduce depth by 70.2% in geomean and cost by 53.8% in geomean.

2 Background

The input to an LSQC problem consists of two parts. The first part is a *quantum program* or a *quantum circuit* which is specified as a list of quantum gates or operations $g_0 g_1 \dots g_{L-1}$ on logical qubits $q_0, q_1, \dots, q_{M-1}$. (There are, in total, L gates and M logical qubits.) Because we assume that logic synthesis is performed prior to LSQC, all of these gates are either single-qubit or two-qubit gates and must be in the gate library of the architecture. An important single-qubit gate is X gate, which negates the qubit; an important two-qubit gate is CX gate, which does not modify the first qubit but changes the second qubit to the XOR of the first and second qubits. Until further notice, we use the Qiskit gate library, which contains X and CX mentioned above, and also other single-qubit gates such as H , T , $T^\dagger$, and S . The specific functions of these gates do not concern LSQC; in fact, we only care about on which qubit(s) the gates operate.

For example, the beginning of a quantum adder program introduced in [2] is shown below.

```
OPENQASM 2.0;           // code for quantum adder
include "qelib1.inc";    // gate library
qreg q[4];              // declare 4 logical qubits
x q[0];                 // g_0
x q[1];                 // g_1
h q[3];                 // g_2
```

Figure 1. Circuit diagram for quantum adder

Figure 2. Coupling Graphs

```
cx q[2], q[3];          // g_3
t q[0];                 // g_4
...                     ...
```

In the above program, g_0 is an X gate on logical qubit q_0 ; g_3 is a CX gate, on q_2 and q_3 . One way to visualize a quantum program is a *circuit diagram*, e.g., Figure 1, where each wire stands for a logical qubit, and the gates are placed from left to right on the wires according to the program order. To make a distinction between single-qubit and two-qubit gates, we separate them into two lists G_1 and G_2 . For $g_l \in G_1$, we use $g_l.q$ to denote the logical qubit it operates on; for $g_l \in G_2$, we use $g_l.q$ and $g_l.q'$. For example, $g_0.q = q_0$, $g_5.q = q_1$, $g_3.q = q_2$, $g_3.q' = q_3$.

The second part of the input to LSQC problem is a coupling graph (P, E) , where $P = \{p_0, p_1, \dots, p_{N-1}\}$ is the set of physical qubits and $E = \{e_0, e_1, \dots, e_{K-1}\}$ is the set of (undirected) connections between them. (There are, in total, N physical qubits and K edges.) The coupling graphs used in this paper are shown in Figure 2. We denote an edge as $(e_k.p, e_k.p')$, e.g., in Figure 2a, $e_1.p = p_0$ and $e_1.p' = p_2$. In addition, fidelity information can be provided as three functions: $f_0 : P \rightarrow [0, 1]$ for measurements, $f_1 : P \rightarrow [0, 1]$ for single-qubit gates, and $f_2 : E \rightarrow [0, 1]$ for two-qubit gates. (Under current technology constraints, quantum computer providers usually only offer one type of entangling two-qubit gate. We are referring the fidelity for this type of gate.)

The output of LSQC is the *spacetime coordinates* (t_l, x_l) for all the input gates and the SWAP gates inserted, and a final mapping $\pi : Q \rightarrow P$ providing which physical qubit to measure for each logical qubit. We show the LSQC result of the quantum adder (Figure 1) on IBM QX2 (Figure 2a) in Figure 3. In this diagram, gates are aligned according to their time coordinates, e.g., $t_0 = t_1 = t_2 = 0$ and $t_{10} = 8$; the space coordinates of single-qubit gates can be read off from the mapping displayed above the wires, e.g., q_0 is mapped

Figure 3. Result of LSQC for quantum adder

Figure 4. Immediate dependencies in the quantum adder (Red arrows are used in OLSQ. Green arrows are those imposed implicitly by the approach of Wille et al.’s. Brown means intersection.)

to p_3 at time 0, so $x_0 = p_3$; the space coordinates of two-qubit gates can be deduced from the mapping, e.g., at time 8, q_3 and q_0 are mapped correspondingly to p_1 and p_2 , so $x_{10} = e_2$ because e_2 connects p_1 and p_2 . The mapping remains the same as the previous time slot if it is not displayed, e.g., q_1 is still mapped to p_2 at time 1, so $x_5 = p_2$. Thus, the final mapping is just the last mapping displayed, $\pi(q_0) = p_2$, $\pi(q_1) = p_3$, $\pi(q_2) = p_0$, and $\pi(q_3) = p_1$. A SWAP gate consisting of three CX gates is inserted on e_3 between p_2 and p_3 . We use the last time slot each SWAP gate takes as its time coordinate (in this case, 7).

The inserted SWAP gate in this example is absolutely necessary. The quantum program has two-qubit gates between q_0 and q_1 (g_8 , g_{12} , and g_{18}), q_1 and q_2 (g_{11}), q_2 and q_3 (g_3 , g_9 , g_{13} , and g_{19}), and finally q_3 and q_0 (g_{10} and g_{21}). This means, without any SWAP gates, the logical qubits must be mapped to a set of physical qubits connected like a square. However, the coupling graph, Figure 2a, does not contain such a structure.

We stress a few constraints for LSQC. 1) We assume any kind of gate cancellation/optimization has been performed prior to LSQC, so all the input gates should be executed. 2) The two-qubit gates should be mapped to valid edges on the coupling graph. 3) Avoid collisions: if two gates g_l and $g_{l'}$ act on the same qubit, then they cannot be executed at the same time, i.e., $t_l \neq t_{l'}$. For example, $t_5 \neq t_8$ because both g_5 and g_8 act on q_1 . 4) Respect dependencies: if we have a dependency between g_l and $g_{l'}$, $l < l'$, then their relative order should not change, i.e., $t_l < t_{l'}$. If there are no gates between these two gates, it is called an *immediate dependency*. However, deriving the dependencies in a generic quantum program turns out to be a nontrivial task. By default, we treat all the collision gate pairs above as dependencies. This is a common practice, e.g., also in Murali et al. [15], and also guarantees correctness. Immediate dependencies of the quantum adder is shown as red and brown arrows in Figure 4. With more knowledge on dependencies, we can improve the LSQC solutions, as we shall see in Section 5.3.

There can be multiple objectives for LSQC: 1) *Depth*: the maximal time coordinate of all gates. Due to the limitation of current QC

technology, physical qubits can only function well up to a short ‘lifetime’. Minimizing depth is thus very important. 2) *SWAP Cost or gate count*: the total number of additional SWAP gates. 3) *Fidelity*: the product of fidelity of all the single-qubit gates, two-qubit gates, and measurements. The current quantum computers are error-prone, so maximal fidelity is desired.

3 Related Works

Exact Solutions: There have been multiple previous exact or optimal approaches related to LSQC. Shafaei et al. [19] first finds a qubit mapping on a 2D grid that minimizes expected distances between qubit pairs required by two-qubit gates. Then, if some two-qubit gate still acts on non-adjacent qubits, the xy routing algorithm is applied to move them together. Wille et al. [25, 27] use pseudo-Boolean optimizer and SMT solver to minimize additional gate cost; Siraichi et al. [21] provide a dynamic programming approach to this task. They split the quantum program into individual gates and consider changes to the mapping at every interval. Instead of individual gates, Bhattacharjee et al. [6] splits the program into ‘levels’ of gates that can be executed in parallel. Afterwards, it optimizes depth with integer linear programming to decide the initial mapping and adjustments of mapping between the levels.

We find that either the gate-by-gate or the level-by-level arrangement imposes some additional constraints than arranging the gates with dependency. We shall illustrate this point with the aforementioned LSQC instance. In this case, Wille et al. return a solution with two SWAP gates, but one SWAP would be enough as shown in Figure 3. This is because when arranging gate-by-gate, there is an implicit dependency between each two-qubit gate and the one after it, which are shown as blue and brown arrows in Figure 4. Specifically, there is a dependency from g_{10} to g_{11} , which means that in their result $t_{10} < t_{11}$, but in the optimal solution, Figure 3, $t_{10} > t_{11}$. In fact, When we manually impose an additional dependency from g_{10} to g_{11} onto OLSQ-SWAP (to be specified in the next section), it gives a solution with two SWAP gates as well. In essence, the gate list is only one of the topological orderings of the dependency graph like Figure 4. If we use a gate-by-gate arrangement with the input gate list order, the result certainly respects all the dependencies, but there are some additional arbitrariness that may lead to sub-optimality. OLSQ actually reconstructs the dependency graph from the gate list and thus effectively explores *all* topological orderings. As for arranging level-by-level, in the LSQC instance above, suppose we assign g_8 and g_9 as the first level, g_{10} and g_{11} as the second level, g_{12} and g_{13} as the third level. Then there must be at least one SWAP inserted between level one and two; otherwise there exists a mapping that can satisfy all the two-qubit connections required by g_8 , g_9 , g_{10} , and g_{11} . This is to say the coupling graph contains a square-like structure, which is impossible in Figure 2a. Similarly, there must be at least one SWAP inserted between level two and three. Thus, at least two SWAP gates are required, which is sub-optimal compared to Figure 3.

Heuristic Solutions: Maslov et al. [14] publish one of the earliest work in LSQC and propose a heuristic algorithm that recursively cuts the coupling graph to derive SWAP gates to adjust the mapping. Siraichi et al. [21] find the initial mapping by matching out-degrees and search for SWAP gates to apply based on a heuristic distance between qubits. Zulehner et al. [28] partition the program and employs A* search with a lookahead scheme to find the SWAP gates

that transform the current mapping to the next. Li et al. [13] apply bi-directional search to minimize additional gate cost. Childs et al. [8] extend existing approaches to the token swapping problem in order for SWAP insertion. Tannu&Qureshi [24] observe the variation in gate fidelity and thus propose taking fidelity into the LSQC problem. Murali et al. [15] present an end-to-end LSQC flow for fidelity optimization. An SMT solver is applied to derive the initial mapping. For SWAP insertion, they construct a reliability matrix of two-qubit gates, which considers the SWAP gates required along the way, and use this matrix to route qubits on the coupling graph. In t|ket> [22], when some two-qubit gate is not executable, the algorithm keeps the relevant qubit permutations as candidates and treats the gates ahead to evaluate them. If more than one candidate receives the highest score before the next non-executable gate, the process goes into a new level until one candidate wins. Recently, Alam et al. [1] design a compilation tool specifically for QAOA circuits. They first assign gates to layers, but due to the commutation relations, most of the dependencies can be ignored, so they apply a two-level search: at the higher level, the order of the layers is decided, followed by a lower level breadth-first search to find the SWAP gates.

Benchmarks: To evaluate these approaches to LSQC, benchmarks such as reversible functions in RevLib [26], gate optimization results of useful logic functions in Amy et al. [2] and Nam et al. [16], or certain circuits with known optimal, QUEKO [23], can be used. The QUEKO evaluations show t|ket> [22] to be a leading industry LSQC tool, but still with large optimality gaps. This serves as a strong motivation for more research into LSQC tools.

4 Approach

In this section, we discuss preprocessing, the objectives, the variable encoding scheme, and the constraints of our proposed optimal layout synthesizer for quantum computing (OLSQ). Then, we introduce some variations to the notion of time to make the synthesizer transition-based (TB-OLSQ), which greatly increases efficiency with little or no performance degradation. Lastly, we consider commutation to improve TB-OLSQ for QAOA circuits.

4.1 Preprocessing

From the input program, we derive a *collision list* C : if two gates g_l and $g_{l'}$, $l < l'$, act on a same logical qubit, then we append $(g_l, g_{l'})$ to C . By default, we use the collision list as the *dependency list* D . Users are also welcome to input their own D based on knowledge of the program. With the dependency list, we can also derive the longest dependency chain with $O(L^2)$ time. This serves as a lower bound of depth to the LSQC result because the dependency chain can only be lengthened by the SWAP gates and cannot be shortened in any case. We will need a time coordinate upper bound T in the formulation. In the hope of a depth-optimal result, we use the longest dependency chain length as T in the beginning.

We also need to extract some features of the coupling graph. We compute an *overlapping edge pair set* O : $\forall e, e' \in E$, $e' \neq e$, if e and e' share some node, append the pair (e, e') to O . We also compute an edge set E_p for each node p : $\forall e \in E$, if $e = (\cdot, p)$ or $e = (p, \cdot)$, we append e to E_p . It is straightforward that $E_p \subset E$ and $\cup_{p \in P} E_p = E$.

4.2 Encoding Variables

- Mapping π_q^t : at time t , logical qubit q is mapped to the physical qubit π_q^t , $\pi_q^t \in P$.

- Time coordinates t_l : gate g_l is being executed at time t_l , $0 \leq t_l \leq T - 1$.
- Space coordinates x_l : if $g_l \in G_1$, then logical qubit $g_l.q$ is mapped to physical qubit x_l , $x_l \in P$; if $g_l \in G_2$, then the two physical qubits, to which $g_l.q$ and $g_l.q'$ are mapped, are connected by edge x_l , $x_l \in E$.
- Use of SWAP gate σ_k^t : if $\sigma_k^t = 1$, then there is a SWAP gate on edge e_k and the last time slot it takes is t (as SWAP gates may take multiple time slots); otherwise, $\sigma_k^t = 0$.

4.3 Constraints

Note that we differentiate variable assignment = and comparison ==. The latter returns true if and only if the equality holds. There is an additional parameter S in our model which stands for the number of time slots a SWAP gate requires. S can be set according to different architectures. We set $S = 3$ as default, as seen in Figure 3.

4.3.1 Injective Mapping. Different logical qubits should be mapped to different physical qubits at any specific time

$$\pi_q^t \neq \pi_{q'}^t \quad \text{for } 0 \leq t \leq T - 1, q, q' \in Q \text{ and } q \neq q' \quad (1)$$

4.3.2 Avoiding Collisions and Respecting Dependencies.

$$t_l < t_{l'} \quad \text{for } (g_l, g_{l'}) \in D \quad (2)$$

4.3.3 Consistency between Mapping and Space Coordinates.

There are two ways we can derive where a gate g_l is at physically: 1) directly through its space coordinate x_l ; 2) indirectly from the mapping of the logical qubit(s) it acts on at its time coordinate, i.e., $\pi_{g_l.q}^{t_l}$ for single-qubit gates; these two should be consistent.

$$(t_l == t) \Rightarrow (\pi_{g_l.q}^t == x_l) \quad \text{for } 0 \leq t \leq T - 1, g_l \in G_1 \quad (3)$$

$$\begin{aligned} [(t_l == t) \wedge (x_l == e)] \Rightarrow \\ \left\{ \left[(\pi_{g_l.q}^t == e.p) \wedge (\pi_{g_l.q'}^t == e.p') \right] \vee \right. \\ \left. \left[(\pi_{g_l.q}^t == e.p') \wedge (\pi_{g_l.q'}^t == e.p) \right] \right\} \\ \text{for } 0 \leq t \leq T - 1, g_l \in G_2, e \in E \end{aligned} \quad (4)$$

4.3.4 Proper SWAP Insertion. Since a SWAP gate takes S time slots, before time $S - 1$, no SWAP gates can finish:

$$\sigma_k^t = 0 \quad \text{for } 0 \leq t \leq S - 2, 0 \leq k \leq K - 1 \quad (5)$$

A SWAP gate cannot overlap with other SWAP gates on the same edge:

$$\begin{aligned} (\sigma_k^t == 1) \Rightarrow (\sigma_{k'}^{t'} == 0) \quad \text{for } S - 1 \leq t \leq T - 1, \\ t - S + 1 \leq t' \leq t - 1, 0 \leq k \leq K - 1 \end{aligned} \quad (6)$$

If two edges overlap in space, the SWAP gates on them cannot overlap in time:

$$\begin{aligned} (\sigma_k^t == 1) \Rightarrow (\sigma_{k'}^{t'} == 0) \quad \text{for } S - 1 \leq t \leq T - 1, \\ t - S + 1 \leq t' \leq t, (e_k, e_{k'}) \in O \end{aligned} \quad (7)$$

A SWAP gate should not overlap with any input single-qubit gates at any time:

$$\begin{aligned} \{(t_l == t') \wedge [(x_l == e_k.p) \vee (x_l == e_k.p')]\} \Rightarrow (\sigma_k^t == 0) \\ \text{for } S - 1 \leq t \leq T - 1, t - S + 1 \leq t' \leq t, 0 \leq k \leq K - 1, g_l \in G_1 \end{aligned} \quad (8)$$

A SWAP gate on e_k should not overlap with any input two-qubit gates on the same edge or the edges that overlap at any time:

$$\left[(t_l == t') \wedge (x_l == e_{k'}) \right] \Rightarrow (\sigma_k^t == 0) \quad \text{for } g_l \in G_2, \quad (9)$$

$$S - 1 \leq t \leq T, \quad t - S + 1 \leq t' \leq t, \quad (e_k, e_{k'}) \in O \text{ or } k' = k$$

4.3.5 Mapping Transformations by SWAP Gates. Mapping at the next time slot is the same with the current one if there are no SWAP gates finished on all the edges in the edge set E_P :

$$\left[\left(\pi_q^t == p \right) \wedge \left(\bigwedge_{e_k \in E_P} \sigma_k^t == 0 \right) \right] \Rightarrow (\pi_q^{t+1} == p) \quad (10)$$

for $0 \leq t \leq T - 2, p \in P, q \in Q$

If there is a SWAP gate finished at t , there can only be one. (Otherwise, the two SWAP gates are on two edges that overlap. This case would be ruled out by Equation 7.) The mapping at $t + 1$ is then transformed by the SWAP gate:

$$\left[\left(\pi_q^t == e_k \cdot p \right) \wedge \left(\sigma_k^t == 1 \right) \right] \Rightarrow (\pi_q^{t+1} == e_k \cdot p')$$

$$\left[\left(\pi_q^t == e_k \cdot p' \right) \wedge \left(\sigma_k^t == 1 \right) \right] \Rightarrow (\pi_q^{t+1} == e_k \cdot p) \quad (11)$$

for $0 \leq t \leq T - 2, 0 \leq k \leq K - 1, q \in Q$

4.4 Objectives

With the set of variables defined above, it is easy to construct the common objectives. In fact, as long as a quantity can be defined from the above variables, it can be the objective. **1) Depth:** $d := \max_{0,1,\dots,L-1} t_l$. We do not need to consider the time coordinates of the inserted SWAP gates because, if a SWAP gate has even larger time coordinate than d defined above, it finishes after all the input gates, thus has no effects on the program and should be ignored. **2) SWAP cost:** $c := \sum_{k=0}^{K-1} \sum_{t=0}^{T-1} \sigma_k^t$. **3) (log-) Fidelity:**

$$f := \sum_{p \in P} \log f_0(p) \left[\sum_{q \in Q} (\pi_q^{T-1} == p) \right] + \sum_{p \in P} \log f_1(p) \left[\sum_{g_l \in G_1} (x_l == p) \right]$$

$$+ \sum_{e \in E} \log f_2(e) \left[\sum_{g_l \in G_2} (x_l == e) \right] + \sum_{k=0}^{K-1} \log f_{\text{SWAP}}(e_k) \left[\sum_{t=0}^{T-1} \sigma_k^t \right] \quad (12)$$

f_0, f_1 , and f_2 above are given as input to the LSQC problem. $f_{\text{SWAP}}(e)$ is the fidelity of a SWAP gate on edge e , which should be computed from the provided single-qubit and two-qubit gate fidelity, depending on how SWAP gates are implemented on the specific architecture. In our case, a SWAP gate consists of three CX gates, so $\log f_{\text{SWAP}} = 3 \log f_2$. To use addition rather than multiplication in the objective, we take the log of f_0, f_1 , and f_2 . To be compatible with other data and variables, which are all integer, we scale up every log fidelity value by 1000 and round it to the nearest integer.

4.5 Complexity Analysis

In our formulation, there are in total $MT + 2L + KT + 2N + 2K$ variables. For regular planar graphs, which most coupling graphs are, the number of edges is usually asymptotically linear to the number of nodes. For example, in a grid, each edge connects two nodes and each node spans out four edges, so $K/N \approx 2$. Therefore, the total number of variables in our formulation is $O(NT)$ where N is the physical qubit count and T is the time coordinate upper bound.

Table 1. Complexity of OLSQ and Related Works

	Solver	Variables	Constraints
[25]	SMT	$O(L_2 \cdot N!)$	$O(L_2 \cdot M \cdot N!)$
[6]	ILP	$O(T \cdot (N^4 + I))$	$O(T \cdot (N^4 + I^2 + N \cdot L))$
OLSQ	SMT	$O(T \cdot N + L)$	$O(T \cdot N \cdot L)$
TB-OLSQ	SMT	$O(B \cdot N + L)$	$O(B \cdot N \cdot L)$

Note: L_2 is the two-qubit gate count. N is the physical qubit count. T is the time coordinate upper bound. I is the number of levels of gates. L is the total gate count. B is the number of gate blocks.

Figure 5. Quantum adder in transition-based model (Two ‘x’s connected by a vertical line segment represent a SWAP gate.)

bound. The total search space is then exponential to N and T . This is expected from the NP-completeness of the problem [21, 23]. However, as shown in Table 1, this formulation still has exponentially fewer variables compared to Wille et al. [25] because they have a variable for each permutation of qubits at each time. OLSQ also has polynomially fewer variables than Bhattacharjee et al. [6], because they also require variables encoding the mapping from *pairs* of logical qubits to pairs of physical qubits at each time. Also, it is straightforward from Table 1 that our formulation reduces the number of constraints significantly.

4.6 Optimality of OLSQ

After we passed the variables, one of the objectives, and the constraints to Z3 SMT solver [9], it would either return a model containing all the variable values that optimizes the given objective, or return ‘unsatisfiable’. As mentioned before, we initially set the time coordinate upper bound T to the largest length of dependency chain. However, it may be the case that on the given architecture, it is impossible to find a solution with this upper bound, e.g. LSQC of the quantum adder, shown in Figure 1, on IBM QX2 coupling graph, shown in Figure 2a. Thus, if the model is unsatisfiable, we geometrically increase T each time by $(1 + \epsilon)x$ until it is satisfiable. We set $\epsilon = 0.3$ in our experiment. This means that OLSQ is optimal up to a certain time coordinate upper bound T . For depth optimization, the optimality is guaranteed. However, for SWAP cost and fidelity optimization, sometimes increasing T even more can lead to better results. However, this is a very rare case as we shall see in Section 5.1, especially when the longest dependency chain in the input quantum program is already of considerable length.

4.7 TB-OLSQ: Rethinking Time Coordinates

In the example of quantum adder, the time upper bound is $T = 15$. However, the mapping to physical qubits changed only once at time 8. Thus, for any specific logical qubit q , the variables π_q^t are the same from $t = 0$ to 7 and $t = 8$ to 14. This is a huge redundancy. In addition, the total search space for the solver is exponential to N and T . Although a cutting edge quantum processor has only

$N = 53$ [3], T as determined by quantum programs can easily grow to be quite large.

These two observations motivate us to improve efficiency of OLSQ by rethinking time coordinates. Instead of keeping π_q^t for all t , it turns out that we can keep only these variables between two *transitions* of the mapping. Formally, a transition is a set of parallel SWAP gates. In the quantum adder example, there is only one transition and the transition consists of only one SWAP gate on edge $e_3 = (p_2, p_3)$. SWAP gates on overlapping edges can not be in parallel, according to Equation 7, so e_1, e_2, e_4 , and e_5 cannot be in the same transition with e_3 . However, e_0 and e_3 together is a valid transition.

Now, we consider a new model of execution which separates the input gates and the inserted SWAP gates: executing some input gates, then a few SWAP gates to make transition(s) in mapping, execute some more input gates, and make other transition(s), ... We can have consecutive transitions without executing any input gates in between. This model is similar to the one in Wille et al. [25], but gates later in the input can appear at the front as long as permitted by dependency, which they do not allow. The quantum adder example in this model is shown in Figure 5, where we first execute the gates in gate block 0, then a SWAP gate to make transition 0, finally the gates in gate block 1. In this transition-based model, there is no notion of precise time. Instead, all the gates in a gate block are at a same coarse-grain time slot, and share the same mapping. This way, the number of mapping variables greatly decreases. In this particular example, there are only 8 mapping variables compared to the original 60 mapping variables.

With slight changes of formulation, OLSQ can be made transition-based. We change the $<$ in Equation 2 to $\leq$, since now even if two input gates have dependency, they can still be assigned to the same gate block, meaning their coarse-grain time coordinates can still be the same. We set $S = 1$ and remove Equation 8 and Equation 9. SWAP gates are separated from input gates in the current model, so we do not need to consider overlaps of input gates with them. $S = 1$ because we are using a coarse-grain time model.

The coarse-grain time upper bound T is initially set to 1, so the solver will search for a solution without any transition. If the solver returns ‘unsatisfiable’, we will increase T by 1 each time until it finds a solution that optimizes the given objective. The value of depth would just be $T - 1$, since there are exactly $T - 1$ gate blocks in the resulting circuit. If SWAP cost or fidelity is set as objective, TB-OLSQ will find the optimal solution that has up to $T - 1$ transitions. Just like OLSQ, there may be better solutions if T is increased even more. After a solution such as Figure 5 is returned, we can use as-soon-as-possible (ASAP) scheduling to derive all the exact time coordinates of the gates. After scheduling, the resulting format is exactly the same as that of OLSQ. We shall show that TB-OLSQ produced optimal or near-optimal solution with orders-of-magnitude speedup compared to OLSQ.

4.8 QAOA-OLSQ: Removing False Dependencies

One of the most important concepts in quantum mechanics is commutation. Only when two operators commute, can we simultaneously ‘observe’ them and receive definite values. The famous Heisenberg’s uncertainty principle states that position and momentum of a particle do not commute, so we cannot simultaneously receive definite values of these two. In the case of quantum computing, if two gates g_l and $g_{l'}$ ($l < l'$) commute, we can change

Figure 6

their relative order without altering the whole program. This means that even if they consecutively act on the same qubit q_m , t_l is not necessarily smaller than $t_{l'}$. Since commutation is purely logical and has nothing to do with the architecture, one may think that it is solely the job of logic synthesis to experiment with the commutation relations. Like many other related works, both OLSQ and TB-OLSQ assume that any commutation is performed prior to LSQC and thus will indeed add $(g_l, g_{l'})$ as a dependency, eliminating the possibility of $t_l > t_{l'}$. However, it turns out that more knowledge of dependencies in LSQC is very beneficial — especially on the QAOA programs [10, 12], which shows great promise to solve some approximate discrete optimization problems.

The QAOA program first sets the qubits to the equal superposition state by applying Hadamard H gates on all the qubits. Then it goes into many iterations. Each iteration has two stages: phase separation and mixing. A simple QAOA program with only one iteration is shown in Figure 6a [11]. Phase separation is implemented by a few ZZ gates, which are two-qubit gates with a parameter γ ; mixing is implemented by R_x gates on all the qubits. The specific functions of these gates are not of concern to Layout synthesis; what matters is on which qubit(s) they act. Since single-qubit gates are always executable, the mixing stage does not require layout synthesis. However, phase separation may contain a ZZ gate on any pair of qubits, which means that layout synthesis is required to move the non-adjacent qubits together when a ZZ gate needs to act on them.

If we input the QAOA program, Figure 6a, and IBM QX2 coupling graph, Figure 2a, to OLSQ or TB-OLSQ, the best result is probably just an identity mapping and the output looks the same with Figure 6a. OLSQ and TB-OLSQ cannot reduce depth because the four ZZ gates all depend on the gate before it, imposing the default dependencies. The particularity of QAOA is that all the ZZ gates commute, which means even both $ZZ(\gamma)(q_2, q_3)$ and $ZZ(\gamma)(q_3, q_4)$ act on q_3 , the latter can actually commute ‘through’ the former. As a result, we have a better phase separation subroutine, as shown in Figure 6b, which has smaller depth. Current qubits still have short ‘lifetimes’ and many QAOA applications have large numbers of iterations, so reducing depth is crucial.

We can improve TB-OLSQ for the phase separation stage of QAOA with the knowledge of commutation. Previously, we treat every collision as a dependency. However, in the phase separation stage, all the ZZ gates are commutable, so none of the collisions are real dependencies. Thus, we simply remove the constraints in Equation 2 in TB-OLSQ. Up to this point, the result would be blocks of original ZZ gates with the fewest transitions possible to make all the qubit pairs adjacent required by these ZZ gates. Inside the ZZ gate blocks, there may be further opportunities to reduce depth with the help of commutation. Therefore, we input this result to OLSQ with depth as objective and, again, remove the constraints in Equation 2. Since the gates are already mapped to valid edges on

Note: The devices with 5, 6, 8, and 16 qubits are IBM QX2, a 2 by 3 grid, a 2 by 4 grid, and Rigetti Aspen-4. Dashed line is an exponential fit of Wille et al.'s results.

Figure 7. Runtime Scaling of Wille et al., and OLSQ-SWAP

the coupling graph, OLSQ does not need to insert any new SWAP gates. Thus, we disable all the σ_k^z variables in OLSQ for speedup. In the end, we derive the spacetime coordinates of all the ZZ gates and the SWAP gates inserted. The depth of this result is highly optimized by the two passes of TB-OLSQ and OLSQ.

5 Evaluation

The evaluations were run on an Ubuntu 16.04 server with two Intel Xeon E5-2699v3 CPUs and 128GB memory. Wille et al. [25]² and TriQ [15]³ were built with Cmake 3.13.4 and GNU Make 4.1. The version of Python was 3.8.2. The versions of Python packages used were Cirq 0.8.0, Pytket 0.5.4, Qiskit 0.18.0, and Z3-Solver 4.8.7.0. We linked the Z3 library contained in Z3-Solver package to Wille et al. and TriQ when building.

We selected a comprehensive set of benchmarks from various sources including [2, 11, 16, 23]. We used the fidelity profile of IBM QX2, Figure 2a, and IBM Melbourne 2d, from [15]. To evaluate fidelity, we inputted the result from different synthesizers to Qiskit, decomposed, and calculated the product of all gate fidelity. We made all the benchmarks and detailed results open-source.⁴

5.1 OLSQ versus Previous Optimal Approaches

The leading exact approach by Wille et al. [25] is open-source. The coupling graph in their original paper is directed, which means that CX gate can only execute in one direction. However, the edges in our formulation is bi-directional. For fair comparison, we input each bi-directional edge as two uni-directional edges to their software. We observe that, compared to the cases where directed coupling graphs are used, their software runs significantly faster. Thus, the runtime data we collect, for comparison with OLSQ, are smaller than those appeared in the original reference [25]. Since Wille et al. aims to minimize SWAP cost, it is most appropriate to compare it against OLSQ with SWAP cost as objective, denoted as OLSQ-SWAP below. We find that, on all instances of benchmark program and architecture, OLSQ-SWAP matches the performance of Wille et al. and sometimes is even better.

Wille et al. have a set of variables denoting whether each qubit permutation is performed between two two-qubit gates. This means that the number of variables in their formulation is proportional to

$N!$, where N is the physical qubit count. This complexity explosion can be observed from the two examples in Figure 7, where the runtime of Wille et al. is nearly exponential. On the other hand, the runtime of OLSQ-SWAP does not show this exponential growth. Wille et al. also rely on pre-processing to derive a function from each qubit permutation to its SWAP cost. Suppose each function value takes 4B memory, then, even for device with qubit count $N = 12$, the memory required is 2TB. As examples, we evaluate three LSQC instances of QUEKO benchmarks and the 16-qubit architecture Aspen-4, Figure 2c. All these evaluations of Wille et al. are aborted because our server runs out of the memory limit we set (32GB). In contrast, OLSQ-SWAP reach the optimal SWAP cost (0) within a relatively short period of time.

5.2 TB-OLSQ versus Heuristic Approaches

One of the leading industry works, t|ket> [22], mainly focuses on optimizing cost. We compare it with TB-OLSQ with SWAP cost set as objective, denoted as TB-OLSQ-SWAP below. One of the leading academic works, TriQ [15], mainly focuses on optimizing fidelity, so we compare it to TB-OLSQ with fidelity set as objective, denoted as TB-OLSQ-Fidelity below. The results are given in Figure 8.

The evaluations denoted as blue dots are the same LSQC instances given to the optimal approaches, so that we can compare heuristic and exact approaches. We find that TB-OLSQ-SWAP has no observable degradation on CX cost compared to OLSQ-SWAP, neither does TB-OLSQ-Fidelity compared to OLSQ-Fidelity on fidelity. This means that our transition-based approach is almost exact; it also hugely increases efficiency, e.g. solving queko_15_1 takes the former 9E4 seconds while taking the latter 30 seconds. Compared to t|ket>, TB-OLSQ-SWAP often reduces the CX cost by large margins, 69.2% in geometric mean. Compared to TriQ, TB-OLSQ-Fidelity often increases the fidelity, even up to 10.0x on some larger programs. In some cases, TriQ has slightly better fidelity results by leveraging larger depths. This is expected since TB-OLSQ-Fidelity optimizes fidelity with the fewest transitions possible, but the number of transitions has no direct link to fidelity, so it is possible to find solutions with higher fidelity and more transitions. However, compared to TriQ, the fidelity loss of OLSQ in these cases remain less than 5%.

5.3 QAOA-OLSQ

Arute et al. [4] is considered a leading work on QAOA implementation, where t|ket> is used to solve LSQC problems of QAOA programs for 3-regular graphs. We conduct evaluation with the same settings in [4]: the coupling graph is part of Google Sycamore with 23 physical qubits, as shown in Figure 2b. We generate random 3-regular graphs with node count M from 10 to 22. (Graph theory shows that the product of M and the vertex degree must be even, so M must be even.) Then, for each edge (i, j) in a 3-regular graph, we append a corresponding gate $ZZ(q_i, q_j)$ to the phase separation stage. The phase separation is given to t|ket>, TB-OLSQ, and QAOA-OLSQ for LSQC and their results are fed to Cirq for statistics on SWAP cost and depth. We consider all ZZ gates and SWAP gates to have unity depth. As shown in Figure 9, even without considering the commutation relations, TB-OLSQ reduces depth by 59.5% in geometric mean and reduces SWAP cost by 29.4% in geometric mean. QAOA-OLSQ further reduces both depth and SWAP cost, by 70.2% and 53.8% in geometric mean compared to t|ket>.

²https://github.com/iic-jku/minimal_ibm_qx_mapping

³<https://github.com/prakashmurali/TriQ>

⁴<https://github.com/UCLA-VAST/OLSQ>

Figure 8. Evaluation of TB-OLSQ

Figure 9. Evaluation of QAOA-OLSQ

6 Conclusion and Future Work

In this paper, we formulate layout synthesis for quantum computing as optimization problems. We present two synthesizers: an exact layout synthesizer (OLSQ) and an approximate, transition-based synthesizer (TB-OLSQ). In addition, we improve TB-OLSQ for QAOA programs (QAOA-OLSQ) by considering commutation. Compared to previous exact approaches, OLSQ is the first that features guaranteed optimality and efficiency both in time and space for general quantum processors. TB-OLSQ shows no visible degradation than OLSQ and is able to outperform leading academic and industry works on multiple metrics. QAOA-OLSQ can reduce both the cost and the depth in a real QAOA experiment setting, which means it would be beneficial for realistic applications for near-term quantum computers. The basic formulation we produce is very versatile in that it can support different objectives, so in the future it can be used even with more complicated objectives involving, say, correlated errors. It is also very valuable to improve the synthesizers based on prior knowledge of either the program,

e.g., commutation relations, or the coupling graph, e.g., special structures inside the graph. In an era when quantum computers still are volatile and error-prone, layout synthesizers prove to be vital for the execution of all quantum programs to minimize the depth (equivalent to coherence time improvement) and cost, and maximize the fidelity.

Acknowledgments

This work is partially supported by NEC under the Center for Domain-Specific Computing Industrial Partnership Program. We would like to thank Iris Cong for valuable comments on the manuscript, the Google Quantum team (Ryan Babbush, Dave Bacon, Craig Gidney, and Matthew Harrigan) for suggesting the QAOA application, Xiaoning Bian, Dmitri Maslov, and Julien Ross for the help on the arithmetic benchmarks, and Prakash Murali for the help on the TriQ. Jason Cong also appreciates multiple discussions with Eugene Ding on the LSQC problem.

References

- [1] Mahabubul Alam, Abdullah Ash-Saki, and Swaroop Ghosh. 2020. An Efficient Circuit Compilation Flow for Quantum Approximate Optimization Algorithm. In *2020 57th ACM/IEEE Design Automation Conference (DAC)*. IEEE, San Francisco, CA, USA, 6.
- [2] M. Amy, D. Maslov, M. Mosca, and M. Roetteler. 2013. A Meet-in-the-Middle Algorithm for Fast Synthesis of Depth-Optimal Quantum Circuits. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 32, 6 (June 2013), 818–830. <https://doi.org/10.1109/TCAD.2013.2244643> arXiv:1206.0758 [quant-ph]
- [3] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhii, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. 2019. Quantum Supremacy Using a Programmable Superconducting Processor. *Nature* 574, 7779 (Oct. 2019), 505–510. <https://doi.org/10.1038/s41586-019-1666-5> arXiv:1910.11333 [quant-ph]
- [4] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Sergio Boixo, Michael Broughton, Bob B. Buckley, David A. Buell, Brian Burkett, Nicholas Bushnell, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Sean Demura, Andrew Dunsworth, Edward Farhi, Austin Fowler, Brooks Foxen, Craig Gidney, Marissa Giustina, Rob Graff, Steve Habegger, Matthew P. Harrigan, Alan Ho, Sabrina Hong, Trent Huang, L. B. Ioffe, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Cody Jones, Dvir Kafri, Kostyantyn Kechedzhii, Julian Kelly, Seon Kim, Paul V. Klimov, Alexander N. Korotkov, Fedor Kostritsa, David Landhuis, Pavel Laptev, Mike Lindmark, Martin Leib, Erik Lucero, Orion Martin, John M. Martinis, Jarrod R. McClean, Matt McEwen, Anthony Megrant, Xiao Mi, Masoud Mohseni, Wojciech Mruczkiewicz, Josh Mutus, Ofer Naaman, Matthew Neeley, Charles Neill, Florian Neukart, Hartmut Neven, Murphy Yuezhen Niu, Thomas E. O'Brien, Bryan O'Gorman, Eric Ostby, Andre Petukhov, Harald Putterman, Chris Quintana, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Andrea Skolik, Vadim Smelyanskiy, Doug Strain, Michael Streif, Kevin J. Sung, Marco Szalay, Amit Vainsencher, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, and Leo Zhou. 2020. Quantum Approximate Optimization of Non-Planar Graph Problems on a Planar Superconducting Processor. (April 2020). arXiv:2004.04197 [quant-ph]
- [5] Alán Aspuru-Guzik, Anthony D. Dutoi, Peter J. Love, and Martin Head-Gordon. 2005. Simulated Quantum Computation of Molecular Energies. *Science* 309, 5741 (2005), 1704–1707. <https://doi.org/10.1126/science.1113479> arXiv:quant-ph/0604193
- [6] Debjyoti Bhattacharjee, Abdullah Ash Saki, Mahabubul Alam, Anupam Chattopadhyay, and Swaroop Ghosh. 2019. MUQUT: Multi-Constraint Quantum Circuit Mapping on NISQ Computers: Invited Paper. In *2019 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*. IEEE, Westminster, CO, USA, 1–7. <https://doi.org/10.1109/ICCAD45719.2019.8942132>
- [7] Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. 2017. Quantum Machine Learning. *Nature* 549, 7671 (Sept. 2017), 195–202. <https://doi.org/10.1038/nature23474>
- [8] Andrew M. Childs, Eddie Schoute, and Cem M. Unsal. 2019. Circuit Transformations for Quantum Architectures. In *14th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2019) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 135)*, Wim van Dam and Laura Mancinska (Eds.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 3:1–3:24. <https://doi.org/10.4230/LIPIcs.TQC.2019.3> arXiv:1902.09102 [quant-ph]
- [9] Leonardo de Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems (Lecture Notes in Computer Science)*, C. R. Ramakrishnan and Jakob Rehof (Eds.). Springer, Berlin, Heidelberg, 337–340. https://doi.org/10.1007/978-3-540-78800-3_24
- [10] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. 2014. A Quantum Approximate Optimization Algorithm. (Nov. 2014). arXiv:1411.4028 [quant-ph]
- [11] Pranav Gokhale, Ali Javadi-Abhari, Nathan Earnest, Yunong Shi, and Frederic T. Chong. 2020. Optimized Quantum Compilation for Near-Term Algorithms with OpenPulse. (April 2020). arXiv:2004.11205 [quant-ph]
- [12] Stuart Hadfield, Zhihui Wang, Bryan O'Gorman, Eleanor Rieffel, Davide Venturelli, and Rupak Biswas. 2019. From the Quantum Approximate Optimization Algorithm to a Quantum Alternating Operator Ansatz. *Algorithms* 12, 2 (Feb. 2019), 34. <https://doi.org/10.3390/a12020034> arXiv:1709.03489 [quant-ph]
- [13] Gushu Li, Yufei Ding, and Yuan Xie. 2019. Tackling the Qubit Mapping Problem for NISQ-Era Quantum Devices. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems - ASPLOS '19*. ACM Press, Providence, RI, USA, 1001–1014. <https://doi.org/10.1145/3297858.3304023> arXiv:1809.02573 [cs.ET]
- [14] Dmitri Maslov, Sean M. Falconer, and Michele Mosca. 2008. Quantum Circuit Placement. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 27, 4 (April 2008), 752–763. <https://doi.org/10.1109/TCAD.2008.917562> arXiv:2002.09783
- [15] Prakash Murali, Norbert Matthias Linke, Margaret Martonosi, Ali Javadi Abhari, Nhung Hong Nguyen, and Cinthia Huerta Alderete. 2019. Full-Stack, Real-System Quantum Computer Studies: Architectural Comparisons and Design Insights. In *Proceedings of the 46th International Symposium on Computer Architecture - ISCA '19*. ACM Press, Phoenix, Arizona, 527–540. <https://doi.org/10.1145/3307650.3322273> arXiv:1905.11349 [quant-ph]
- [16] Yunseong Nam, Neil J. Ross, Yuan Su, Andrew M. Childs, and Dmitri Maslov. 2018. Automated Optimization of Large Quantum Circuits with Continuous Parameters. *npi Quantum Information* 4, 1 (Dec. 2018), 23. <https://doi.org/10.1038/s41534-018-0072-4> arXiv:1710.07345 [quant-ph]
- [17] Michael A Nielsen and Isaac L. Chuang. 2010. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, Cambridge, United Kingdom.
- [18] Eleanor G. Rieffel, Stuart Hadfield, Tad Hogg, Salvatore Mandrà, Jeffrey Marshall, Gianni Mossi, Bryan O'Gorman, Eugeniu Plamadeala, Norm M. Tubman, Davide Venturelli, Walter Vinci, Zhihui Wang, Max Wilson, Filip Wudarski, and Rupak Biswas. 2019. From Ansätze to Z-Gates: A NASA View of Quantum Computing. (May 2019). arXiv:1905.02860 [quant-ph]
- [19] Alireza Shafaei, Mehdi Saeedi, and Massoud Pedram. 2014. Qubit Placement to Minimize Communication Overhead in 2D Quantum Architectures. In *2014 19th Asia and South Pacific Design Automation Conference (ASP-DAC)*. IEEE, Singapore, 495–500. <https://doi.org/10.1109/ASPDAC.2014.6742940>
- [20] P.W. Shor. 1994. Algorithms for Quantum Computation: Discrete Logarithms and Factoring. In *Proceedings 35th Annual Symposium on Foundations of Computer Science*. IEEE Comput. Soc. Press, Santa Fe, NM, USA, 124–134. <https://doi.org/10.1109/SFCS.1994.365700> arXiv:quant-ph/9508027
- [21] Marcos Yukio Siraichi, Vinicius Fernandes dos Santos, Sylvain Collange, and Fernando Magno Quintao Pereira. 2018. Qubit Allocation. In *Proceedings of the 2018 International Symposium on Code Generation and Optimization - CGO 2018*. ACM Press, Vienna, Austria, 113–125. <https://doi.org/10.1145/3168822>
- [22] Seyon Sivarajah, Silas Dilkes, Alexander Cowtan, Will Simmons, Alec Edgington, and Ross Duncan. 2020. t|ket| : A Retargetable Compiler for NISQ Devices. *Quantum Science and Technology* (April 2020). <https://doi.org/10.1088/2058-9565/ab8e92> arXiv:2003.10611 [quant-ph]
- [23] Bochen Tan and Jason Cong. 2020. Optimality Study of Existing Quantum Computing Layout Synthesis Tools. *IEEE Trans. Comput.* (July 2020). <https://doi.org/10.1109/TC.2020.3009140> arXiv:2002.09783 [quant-ph]
- [24] Swamit S. Tannu and Moinuddin K. Qureshi. 2019. Not All Qubits Are Created Equal: A Case for Variability-Aware Policies for NISQ-Era Quantum Computers. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems - ASPLOS '19 (ASPLOS '19)*. Association for Computing Machinery, Providence, RI, USA, 987–999. <https://doi.org/10.1145/3297858.3304007> arXiv:1805.10224 [cs.ET]
- [25] 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 on - DAC '19*. ACM Press, Las Vegas, NV, USA, 1–6. <https://doi.org/10.1145/3316781.3317859> arXiv:1907.02026 [quant-ph]
- [26] 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*. IEEE, Dallas, TX, USA, 220–225. <https://doi.org/10.1109/ISMVL.2008.43>
- [27] Robert Wille, Aaron Lye, and Rolf Drechsler. 2014. Optimal SWAP Gate Insertion for Nearest Neighbor Quantum Circuits. In *2014 19th Asia and South Pacific Design Automation Conference (ASP-DAC)*. IEEE, Singapore, 489–494. <https://doi.org/10.1109/ASPDAC.2014.6742939>
- [28] Alwin Zulehner, Alexandru Paler, and Robert Wille. 2018. Efficient Mapping of Quantum Circuits to the IBM QX Architectures. In *2018 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, Dresden, Germany, 1135–1138. <https://doi.org/10.23919/DATE.2018.8342181> arXiv:1712.04722 [quant-ph]