



Optimal Layout Synthesis for Deep Quantum Circuits on NISQ Processors with 100+ Qubits

Irfansha Shaik  

Department of Computer Science, Aarhus University, Denmark
Kvantify Aps, Copenhagen S, Denmark

Jaco van de Pol  

Department of Computer Science, Aarhus University, Denmark

Abstract

Layout synthesis is mapping a quantum circuit to a quantum processor. SWAP gate insertions are needed for scheduling 2-qubit gates only on connected physical qubits. With the ever-increasing number of qubits in NISQ processors, scalable layout synthesis is of utmost importance. With large optimality gaps observed in heuristic approaches, scalable exact methods are needed. While recent exact and near-optimal approaches scale to moderate circuits, large deep circuits are still out of scope. In this work, we propose a SAT encoding based on parallel plans that apply 1 SWAP and a group of CNOTs at each time step. Using domain-specific information, we maintain optimality in parallel plans while scaling to large and deep circuits. From our results, we show the scalability of our approach which significantly outperforms leading exact and near-optimal approaches (up to 100x). For the first time, we can optimally map several 8, 14, and 16 qubit circuits onto 54, 80, and 127 qubit platforms with up to 17 SWAPs. While adding optimal SWAPs, we also report near-optimal depth in our mapped circuits.

2012 ACM Subject Classification Hardware → Quantum computation; Computing methodologies → Planning for deterministic actions

Keywords and phrases Layout Synthesis, Transpiling, Qubit Mapping and Routing, Quantum Circuits, Propositional Satisfiability, Parallel Plans

Digital Object Identifier 10.4230/LIPIcs.SAT.2024.26

Related Version *Previous Version*: <https://arxiv.org/abs/2403.11598>

Supplementary Material *Software*: <https://github.com/irfansha/Q-Synth> [24]

archived at `swh:1:dir:be31c57364bd541c6b65afac80603e9004cf4008`

Software (Github Tag): <https://github.com/irfansha/Q-Synth/releases/tag/Q-Synth-v2.0-SAT2024> [25], archived at `swh:1:dir:be31c57364bd541c6b65afac80603e9004cf4008`

Funding IFD (Innovation Fund Denmark)

1 Introduction

The Quantum Layout Mapping problem takes as input a quantum circuit (logical design) and a coupling map (connectedness between physical qubits). The result is an “equivalent” quantum circuit mapped to the physical qubits, such that any binary operation only happens on connected qubits. Besides an initial mapping of logical qubits to physical qubits, this also involves the insertion of SWAP gates. Noise is inherent to qubits in Noisy Intermediate-Scale Quantum (NISQ) processors. Additional SWAP gates increase both the 2-qubit gate count and the circuit depth. In the current NISQ era, minimizing error is of utmost importance for any practical quantum computing. The error rate depends on the number of gates, the fidelity of gates, and the depth of the circuit. The Optimal Quantum Layout Synthesis is to synthesize a mapping that optimizes one of the above metrics.

Optimal Layout Synthesis has been studied before. A nice overview is provided in [29]. Several heuristic approaches exist which optimize various metrics. The classical algorithm for heuristic mapping is SABRE (in Qiskit) [15]. [27] use the MQT benchmarks for mapping



© Irfansha Shaik and Jaco van de Pol;

licensed under Creative Commons License CC-BY 4.0

27th International Conference on Theory and Applications of Satisfiability Testing (SAT 2024).

Editors: Supratik Chakraborty and Jie-Hong Roland Jiang; Article No. 26; pp. 26:1–26:18

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

and swapping, using a heuristic search space reduction with an $O(n \log n)$ algorithm. Other approaches used include A* with cost metrics [32], MAXSAT [19], temporal planning [30], and constraint programming [3] (minimizing circuit depth).

While heuristic approaches are fast and scalable, their suboptimal mappings may result in high error rates [31, 29]. Optimizing fidelity with exact approaches can result in circuits with the lowest error rate. However, as shown in [28], optimizing fidelity is extremely hard and does not scale beyond small circuits. Circuit depth and 2-qubit gate count optimization are better alternatives for scalability. The OLSQ tool¹ optimizes circuit depth and is built on [28]. A scalable variant OLSQ2 based on Z3 appeared in [16]. The QMAP tool² optimizes the number of SWAP gates and is based on [33, 31]. The same authors introduced the use of subarchitectures [22]. Other ideas to improve quantum layout use quantum teleportation [10]. In [4], measurements are placed early so qubits can be reused.

In [26], we proposed a tool, Q-Synth v1³, for SWAP gate optimization which outperformed both QMAP and OLSQ tools. Q-Synth v1 reduces optimal quantum layout synthesis to classical planning. For maintaining the optimality of the SWAP gates added, Q-Synth v1 adds exactly 1 CNOT or 1 SWAP gate at each time step. In such an approach, the hardness increases with the plan length i.e., the number of CNOTs + SWAPs. Despite the recent progress in Q-Synth v1 and OLSQ2, deep circuits that require many SWAPs are still out of reach.

Contribution

In this paper, we provide a SAT encoding based on parallel plans with domain-specific information. In particular, at each time step, we map one SWAP gate and a group of CNOT gates. This reduces the make-span, and using domain-specific information we maintain the optimality. We propose two-way constraints for CNOT dependencies for better dependency propagation. In addition, we also provide variations of our encoding with bridges and relaxed dependencies (via commutation). In all variations, we only add provably optimal number of bridges+SWAPs.

For experimental evaluation, we consider two benchmark sets: 1) Standard benchmarks from previous papers; and 2) Deep VQE benchmarks. For comparison, we consider leading near-optimal tool TB-OLSQ2 [16] and heuristic SABRE [15]. For mapping, we consider 4 NISQ processors Melbourne (14 qubits), Sycamore (54 qubits), Rigetti (80 qubits), and Eagle (127 qubits). We propose three experiments: in the first two experiments we map both benchmark sets to the Sycamore, Rigetti, and Eagle platforms. In the first experiment, we compare the number of SWAPs added by all three tools. In the second experiment, we compare SWAP additions and circuit depth of the mapped circuits with TB-OLSQ2. In the third experiment, we compare the effectiveness of bridges and relaxed dependencies in our tool by mapping onto the Melbourne platform. Here we report the additional number of (optimal) SWAPs+bridges.

We demonstrate that our encoding can optimally map deep circuits onto large platforms with up to 127 qubits. Our tool outperforms the leading near-optimal tool TB-OLSQ2 up to 100x while always adding the optimal number of SWAPs. We show that while adding optimal SWAPs, we also report near-optimal depth in the mapped circuits. We also confirm that heuristic approaches like SABRE add too many SWAPs.

¹ OLSQ tool <https://github.com/tbcdebug/OLSQ>

² Munich Quantum Toolkit QMAP <https://github.com/cda-tum/qmap>

³ Q-Synth v1 tool <https://github.com/irfansha/Q-Synth/releases/tag/Q-Synth-v1.0-ICCAD23>

2 Preliminaries

2.1 Layout Synthesis for Quantum Circuits

A quantum circuit consists of a fixed number of (logical) qubits, and a number of quantum gates (operations) that are applied to some qubits in a particular order. If the output qubit of gate g_1 is used as an input qubit of gate g_2 , we say g_2 depends on g_1 . The dependencies form a DAG (directed acyclic graph) between the gates. Gates that are (transitively) independent are called parallel, and can be applied in any order.


Any quantum circuit can be decomposed to an intermediate representation with only single-qubit gates and CNOT gates [7]. Viewed classically, the binary CNOT gate (controlled-NOT, also known as CX) takes two qubits (a, b) as input and transforms them into $(a, a \oplus b)$, i.e., the control qubit a determines whether the data qubit b is negated. We will also use the SWAP gate, which transforms a qubit pair (a, b) into (b, a) . A SWAP gate can be expressed as a sequence of 3 CNOT gates.

The single-qubit gates will be treated as black-boxes in this paper; they are distinguished by their name (X, Z, H, S, T, etc.) but we don't make assumptions on their semantics (except in an extension of our method in Section 3.1). We refer the interested reader to [21] for a detailed introduction to quantum gates and quantum circuits in general.

Most physical quantum platforms have limited connectivity, in which the CNOT operations can only be applied on physical qubits that are neighbors in the so-called coupling graph. Given such a circuit and a coupling graph, Layout Synthesis consists of two phases: Initial Mapping and Qubit Routing. In Initial Mapping, the logical qubits of the given circuit are mapped to some physical qubits of the platform bijectively. In Qubit Routing, the following constraints must be satisfied:

- Every gate must be scheduled in an order that respects all dependencies;
- Every gate must be applied to the correct qubits (taken the mapping into account);
- The 2-qubit CNOT gates can only be mapped on connected physical qubits.

Additional SWAP gates may be required, to swap the values of connected physical qubits to ensure all CNOT gates can be mapped. In this paper, we use gate count as an optimization metric. In Layout Synthesis, the number of single-qubit gates and CNOT gates remain unchanged. Optimal Layout Synthesis is thus minimizing the additional SWAP gates.



■ **Figure 1** 3-qubit Or circuit with 6 CNOT gates.

For example, Figure 1 shows an Or-circuit with 3 logical qubits (horizontal lines $\{l_0, l_1, l_2\}$). The circuit has 11 single-qubit gates (boxes with names) and 6 CNOT gates (the dot indicates the control qubit, while the \oplus indicates the data qubit). Let us suppose we want to map this circuit onto the linear 3-qubit platform as in Figure 2b. Regardless of physical qubit connections, single-qubit gates can always be scheduled. Only the 2-qubit CNOT gates are relevant for our optimal synthesis problem. Thus, we first remove the single-qubit gates and only consider CNOT gates for the mapping. After finding the optimal mapping, the single-qubit gates will be reinserted. Figure 2a shows the CNOT gates in the Or-circuit.




Figure 2 Reduced Or circuit and a 3-qubit linear platform.

In any valid mapping, the dependencies must be respected, for example, gates c_1 and c_3 can only be mapped before and after gate c_2 , respectively. In this example, the dependency graph is a total order, but note that with 4 qubits, parallel CNOT gates are possible, which can be scheduled in any order. One can observe that the connections of the CNOT gates c_1, c_2 and c_3 form a triangle $(l_0, l_1), (l_1, l_2), (l_2, l_0)$. Since the coupling graph does not have a triangle, one cannot map our example circuit to the linear platform. At least two SWAP gates are needed for any valid mapping.

Figure 3 shows such a mapped Or-circuit where l_0, l_1, l_2 are mapped to p_0, p_1, p_2 respectively. Intuitively, the SWAP gates slice the circuit such that the sub-circuits do not have any triangles by CNOT connections.




Figure 3 Mapped Or-circuit with 2 additional SWAPs (optimal).

Finally, single-qubit gates can be inserted back respecting original DAG dependencies. Figure 4 shows the final mapped circuit with optimal SWAP gates. Note that the number of physical qubits can be more than logical qubits. In such cases, one can use so-called *ancillary* qubits to avoid unnecessary swaps. Similar to Q-Synth v1, we allow ancillary swapping i.e., a mapped physical qubit can be swapped with an empty physical qubit.




Figure 4 Final Mapped Or-circuit after inserting back single-qubit gates.

2.2 Optimal Layout Synthesis as Planning

In Q-Synth v1, we encoded optimal layout synthesis as a planning problem in the Planning Domain Definition (PDDL) Specification. As discussed above, the reduced circuit with only CNOT gates is mapped using additional SWAPs. Later single-qubit gates are inserted back

to reconstruct the final mapped circuit. In such a planning problem, either exactly one CNOT gate or one SWAP gate is scheduled at each time step. A plan with the optimal number of actions corresponds to the optimal number of SWAP additions.

Planning as SAT

Given a boolean formula, a Satisfiability (SAT) problem is finding an assignment to the boolean variables that makes it a true formula. A planning problem can be encoded as a bounded reachability problem. Sequential encoding [14] is a standard SAT encoding where each time step encodes a single action. Using a sequential encoding, one can obtain optimal plans by incrementing the plan length by 1. For instance, one could use Q-Synth v1 with Madagascar (a SAT-based planner) to find an optimal mapping. Since the optimal plan length for the example, Or-circuit is 8 (6 CNOTs + 2 SWAPs), the SAT instance has a make-span of 8. As shown in [26], sequential encoding scales well for moderate circuits however deep circuits are still out of reach. It is consensus that a long optimal plan length can severely impact the performance of SAT-based planners.

Parallel Plans

In literature, alternative parallel plan encodings like \forall -step [13] and \exists -step [23] were proposed for scaling heuristic SAT-based planning. The key idea in a parallel plan is to group two or more actions whose preconditions and effects do not conflict. While encodings like \exists -step scale well, the optimality is not guaranteed. For a scalable optimal layout synthesis, one needs a way to group CNOTs while still maintaining the optimality.

2.3 Parallel Plans in Optimal Layout Synthesis

Madagascar implements both \forall -step and \exists -step parallel plans. Directly using Q-Synth v1 with parallel plan encodings in Madagascar does not preserve optimality. In particular, there are three main challenges:

- More than one SWAP gate can be applied at each parallel step;
- Planner needs to find a partial order in each parallel step satisfying dependencies;
- Relaxing CNOT dependencies within a parallel step is not trivial in a PDDL specification.

In this paper, we directly encode Layout Synthesis as a SAT problem to circumvent the encoding challenges in a PDDL specification. In our encoding, we allow exactly one SWAP gate at each parallel step. Thus, the number of parallel steps corresponds to the optimal SWAP additions. Further, we use domain-specific information from Layout Synthesis to relax CNOT dependencies within a parallel step. In particular, we make two observations:

- The qubit mappings do not change between two consecutive SWAP gates;
- Given a set of CNOTs, one can always reconstruct a partial order with DAG dependencies.

We take advantage of the partial order reconstruction and drop dependency constraints within a parallel step. The SAT solver can now choose exactly one SWAP and a group of CNOTs in each parallel step. CNOT gates in different time steps still need to respect the original DAG dependencies. The full SAT encoding is later discussed in detail in Section 3.

For our example, a parallel plan with a make-span of 3 is sufficient (see Figure 3) instead of 8. At time step 0, along with the initial mapping, a group of CNOT gates are also mapped. From time step 1, exactly one SWAP gate and a group of CNOT gates are mapped. Note that the satisfying assignment returned by a SAT solver only specifies that:

- Logical qubits l_0, l_1, l_2 are mapped to p_0, p_1, p_2 , respectively;
- SWAP gates on p_1, p_2 and p_0, p_1 are applied at time steps t_1, t_2 respectively;

- c_1, c_2 gates are applied at t_0 ;
- c_3, c_4, c_5 gates are applied at t_1 ;
- c_6 gate is applied at t_2 ;

In mapped circuit reconstruction, we use the DAG dependencies to order the group of CNOTs in each time step. In the literature, an SMT based encoding is applied in TB-OLSQ(2) [28, 16] which also groups CNOT gates between consecutive SWAP gates. They optimize make-span of their defined problem. However, longer make-span may result in better SWAP count and circuit depth. In our experiments, we indeed observe suboptimal solutions by TB-OLSQ2 in both metrics.

2.4 Incremental SAT Solving

Conflict Driven Clause Learning (CDCL) [18] is a key part of state-of-the-art SAT solving. When solving similar instances, one can reuse the learned clauses. Incremental SAT solving allows solving a SAT instance given an assumption of a partial assignment. Essentially by using different assumptions, multiple instances can be solved while reusing the learned clauses. In problems like planning, one needs to refute up to $k-1$ plan length for optimal plans. By adding assumptions encoding that the goal is reached in the current iteration, one can solve a planning instance incrementally.

3 Two-Way Parallel SAT encoding

In this section, we implement the ideas discussed above. We provide an incremental SAT encoding that applies the idea of parallel plans in Layout Synthesis. Table 1 describes the main variables used in the encoding. Algorithm 1 describes the structure of our encoding. In every time step, a group of CNOTs are applied. From time step 1, each incremental step adds one extra SWAP. We generate a set of variables for CNOT and SWAP constraints at each time step.

■ **Table 1** Encoding variables and descriptions.

Variable	Description
$m_{l,p}^t$	mapping var for logical l and physical p qubits at time step t
mp_p^t	if physical qubit p is mapped to some logical qubit at time step t
$s_{p,p'}^t$	SWAP variable for physical qubits p, p' at time step t
st_p^t	SWAP-touched variable for physical qubit p at time step t
$c_i^t / ac_i^t / dc_i^t$	current/advanced/delayed CNOT var for i th CNOT at time step t
$lp_{l,l'}^t$	logic qubit pair variables for logical qubits l, l' at time step t

In addition to specifying which CNOT gates are chosen in each time step, we also need to specify the CNOT dependencies. We use the DAG generated from the original circuit for computing the CNOT dependencies. We adapt the CNOT dependency constraints by specifying that for every CNOT gate in time step t its predecessors (successors) can be applied at time step t' , where $t' \leq t$ ($t' \geq t$). We use two extra CNOT blocks, advanced and delayed CNOTs, which specify if a CNOT gate is mapped in an earlier or later time step. We call this Two-way SAT encoding to emphasize the bidirectional propagation of CNOT dependencies. In the following paragraphs, we describe the main parts of the Algorithm and provide the constraints.

■ **Algorithm 1** Incremental SAT Solving, starting with $t=0$.

```

1: for all  $l \in [1 \dots nl]$  do
2:   ExactlyOne( $m_{l,1}^t, \dots, m_{l,np}^t$ )
3: for all  $p \in [1 \dots np]$  do
4:   AtmostOne( $m_{1,p}^t, \dots, m_{nl,p}^t$ )
5: MappedPQubits
6: if  $t \neq 0$  then
7:   SwapConstraints and Ancillaries
8:   CNOTConnections and CNOTDependencies
9: Assumptions
10: Solve instance with assumption  $asm^t$ 
11: if Instance not satisfied then
12:   repeat from step 1 with  $t = t + 1$ 

```

Initial Mapping

Let L (P) be a set of logical (physical) qubits in the circuit. Let nl (np) be the number of logical (physical) qubits. In time step 0, we add requirements on the Initial Mapping for logical and physical qubits. Lines 1 to 4 in the Pseudo code add constraints for mapping every logical qubit to a unique physical qubit. We use one-hot encoding for specifying the mapping. We apply ExactlyOne (AtmostOne) constraints for logical (physical) qubit mapping variables. Adding these constraints only at the 0th time step is sufficient for correctness. However, adding these constraints at every time step significantly improved the performance of SAT solvers.

SwapConstraints

From time step 1, we use the same mapping variables for handling SWAPs. Adding a SWAP gate changes the mapping between logical and physical qubits. Let CP be the set of all connected physical qubit pairs. The following constraints must ensure that a SWAP gate is only applied on a connected physical qubit pair. The logical qubits mapped on the physical qubit pair must be swapped in the next time step. The qubit mappings for the untouched physical qubits must be propagated. We define two sets of variables to satisfy such constraints. For choosing a SWAP, we define one SWAP variable $s_{p,p'}$ for each connected physical qubit pair (p, p') . For propagation, we define SWAP-touch variables st_p to specify if the physical qubit p is touched by the SWAP. We specify that 1) Each SWAP variable forces the SWAP-touched physical variables to True; 2) Exactly one of the SWAP variables is set to True; 3) Every SWAP forces exactly two SWAP-touched variables to True. Let S be the set of all SWAP variables. The corresponding boolean constraints are:

$$\bigwedge_{(p,p') \in CP} (s_{p,p'}^t \rightarrow st_p^t \wedge st_{p'}^t) \wedge \text{ExactlyOne}(S^t) \wedge \text{ExactlyTwo}(st_1^t, \dots, st_{np}^t)$$

Based on the chosen SWAP variables, we update the mapping variables. For each SWAP variable $s_{p,p'}$, we swap the p and p' mapped variables from the previous to the current step.

$$\bigwedge_{(p,p') \in CP} \bigwedge_{l \in L} s_{p,p'}^t \rightarrow ((m_{l,p}^{t-1} \leftrightarrow m_{l,p'}^t) \wedge (m_{l,p'}^{t-1} \leftrightarrow m_{l,p}^t))$$

If a SWAP-touch variable is False, we propagate the corresponding mapping variables.

$$\bigwedge_{p \in P} \bigwedge_{l \in L} \neg st_p^t \rightarrow (m_{l,p}^{t-1} \leftrightarrow m_{l,p}^t)$$

MappedQubits and Ancillaries

Using additional qubits, traditionally called *ancillaries*, can reduce the total number of SWAPs needed. An *ancillary SWAP* exchanges a mapped qubit with an unmapped qubit. To specify this, we need to keep track of mapped qubits (including at time step 0). We specify that a physical qubit p is mapped to some logical qubit if and only if its mapped variable mp_p is True.

$$\bigwedge_{p \in P} mp_p^t \leftrightarrow \left(\bigvee_{l \in L} m_{l,p}^t \right)$$

We restrict that at least one of the swapped physical qubits is a mapped qubit. With similar constraints, we also provide an option for only non-ancillary SWAPs.

$$\bigwedge_{(p,p') \in CP} s_{p,p'}^t \rightarrow (mp_p^t \vee mp_{p'}^t)$$

CNOTConnections

Let CL be the set of all connected logical qubit pairs derived from the CNOT connections in the input circuit. We require CNOT gates to be applied only on connected physical qubits. Since CNOT gates must be applied to specific logical qubits, we require that the corresponding logical qubits be mapped to connected physical qubits. First, we specify that logical qubit pair variables are true if and only if the physical qubits they are mapped are connected.

$$\begin{aligned} \bigwedge_{(l,l') \in CL} \left(\bigwedge_{(p,p') \in CP} ((m_{l,p}^t \wedge m_{l',p'}^t) \vee (m_{l,p'}^t \wedge m_{l',p}^t)) \rightarrow lp_{l,l'}^t \wedge \right. \\ \left. \bigwedge_{(p,p') \in \overline{CP}} ((m_{l,p}^t \wedge m_{l',p'}^t) \vee (m_{l,p'}^t \wedge m_{l',p}^t)) \rightarrow \neg lp_{l,l'}^t \right) \end{aligned}$$

Using logical qubit pair variables, we specify that if a CNOT is mapped then its corresponding logical qubits are connected. We define \mathcal{D} as a dictionary of CNOT indices to logical qubit pairs. Let nc be the number of CNOT gates. The corresponding boolean constraint is:

$$\bigwedge_{nc}^{i=1} c_i^t \rightarrow lp_{\mathcal{D}[i]}^t$$

CNOTDependencies

For a correct mapping, we need to respect the DAG dependencies of the CNOT gates in a circuit. As discussed earlier, we use advanced and delayed CNOT blocks in each time step to propagate local information globally. Every CNOT is mapped, advanced, or delayed in all time steps, depending on the status of its predecessors (*pre*) and successors (*suc*) in the dependency DAG. If at time step t a CNOT is:

- Mapped: Its predecessors (successors) are either advanced (delayed) or mapped in the same time step.
- Advanced: 1) It is applied or advanced in $t - 1$; 2) Its predecessors are also advanced in t .
- Delayed: 1) It is either applied or delayed in $t + 1$; 2) Its successors are also delayed in t ; 3) In t , either its logical qubits are not connected or one of its predecessors is delayed.

The corresponding boolean constraints are:

$$\begin{aligned}
& \bigwedge_{\text{nc}}^{i=1} (\text{ExactlyOne}(c_i^t, ac_i^t, dc_i^t) \wedge \\
& \quad \bigwedge_{j \in \text{pre}(i)} c_i^t \rightarrow (ac_j^t \vee c_j^t) \wedge \bigwedge_{j \in \text{suc}(i)} c_i^t \rightarrow (c_j^t \vee dc_j^t) \wedge \\
& \quad ac_i^t \rightarrow (c_i^{t-1} \vee ac_i^{t-1}) \wedge \bigwedge_{j \in \text{pre}(i)} ac_i^t \rightarrow ac_j^t \wedge \\
& \quad dc_i^{t-1} \rightarrow (c_i^t \vee dc_i^t) \wedge \bigwedge_{j \in \text{suc}(i)} dc_i^t \rightarrow dc_j^t \wedge dc_i^t \rightarrow (\neg \text{lp}_{\mathcal{D}[i]}^t \vee \bigvee_{j \in \text{pre}(i)} dc_j^t))
\end{aligned}$$

Assumptions for Incremental solving

We specify that CNOT gates cannot be advanced at time step 0 i.e., $\bigwedge_{\text{nc}}^{i=1} \neg ac_i^0$. For using incremental solving in SAT, we use an assumption variable asm^t . At every time step, if the assumption variable is true then CNOT gates cannot be delayed i.e., $\text{asm}^t \leftrightarrow \bigwedge_{\text{nc}}^{i=1} \neg dc_i^t$. For each time step, we call the SAT solver with the assumption variable asm^t as True.

Encoding Size


Let l be the number of logical qubits, p be the number of physical qubits, p_e be the number of edges in the physical coupling graph, l_e be the number of edges in the logical graph (from CNOT gates), c be the number of CNOTs, and finally, let k be the make-span. The encoding requires $O(k(lp + p_e + l_e + c))$ variables. Usually, the physical coupling graphs are planar, so the variables required is $O(k(lp + l_e + c))$. Note that, we use the sequential counter encoding for exactly-one constraints, so it can add extra $O(p)$ auxiliary variables. In total, the number of variables is $O(k(lp + l_e + c))$. The encoding requires $O(k(lp + lp_e + l_e p^2 + c))$ clauses. Again, for a planar physical coupling graph this is bounded by $O(k(lp + l_e p^2 + c))$ clauses.

3.1 Additional Functionality

So far, we have not used the semantics of the unary or binary gates (except the SWAP gates). The previous encoding could also be used to map circuits with for instance binary CZ gates instead of CNOT gates. If we take the semantics of the gates into account, there are more opportunities to optimize the circuits. While a complete re-synthesis of the circuit is beyond the scope of this paper, we want to illustrate some known techniques that can further reduce the number of SWAP gates. We emphasize that we now change the optimization problem (by allowing more solutions), and that the extensions are specific for CNOT gates. Our main purpose is to show that our proposed encoding can be easily extended to incorporate these techniques, known as “bridges” and “relaxed dependencies”. We also note that the encoding can be easily restricted to disallow the use of “ancillary qubits”.

Bridges


Using a bridge, one can apply a CNOT on disconnected physical qubits. For instance, the bridge b_1 - b_4 in Fig. 5 together implements a CNOT between p_0 and p_2 (which implements c_4 after the preceding SWAP). We limit ourselves to bridges of distance 2. Observe that the bridge introduces 3 extra CNOT gates, so the cost is the same as a SWAP gate. However, the result is different, since a bridge does not swap the qubits. This might be an advantage, depending on the rest of the circuit. In [11], it was shown that using bridges can reduce the overall CNOT count.



■ **Figure 5** Mapping reduced circuit using a SWAP and a bridge gate.

We adapted our optimal Two-Way SAT encoding, by allowing to add either a single bridge or a single SWAP gate at each time step. If a bridge was added, the corresponding CNOT gate is regarded as scheduled. So both options cost exactly 3 CNOT gates. The SAT solver will find a solution with the minimal sum of bridge or SWAP gates. Our experiments will show that we indeed find better solutions with bridges.

Relaxed Dependencies



■ **Figure 6** Mapping reduced circuit using a SWAP and relaxed dependencies.

The authors of [12] consider gate commutation rules for quantum layout mapping. Commutation and cancellation rules on R_Z and CNOT gates are also used in [20], to reduce the number of H-gates. For instance, two subsequent CNOT gates on the same control qubit, or on the same data qubit, can be commuted without changing the semantics. Also, single Z-like gates (like the Z-, S-, T- and R_Z -gates) commute with the control bit of a CNOT, while X-like gates (like the X- and R_X -gates) commute with its data bit.

This added freedom can be exploited: by permuting the CNOT gates, they can be grouped in a convenient manner, so less SWAP gates are needed. For instance, in Fig. 6a, c_3 and c_4 can be commuted, since they share their control bit. As a result, we can now find a solution that requires only 1 SWAP gate (Fig. 6b), while still respecting the linear coupling graph.

The gate commutation rules can be incorporated in our optimal Two-way SAT encoding by computing a “relaxed” dependency graph. In the example above, we consider c_2 as a dependency for c_3 and c_4 , but c_3 and c_4 are considered independent. We stress that the relaxed dependencies must also take the unary gates into account (for instance c_2 and c_4 cannot be commuted, since there is a T-gate in between, cf. Fig. 1).

In our tool, we compute the relaxed dependency graph before removing the unary gates. We then generate the encoding as presented before, replacing the dependency graph with the relaxed dependency graph. This guarantees an optimal Layout Mapping (minimal number of SWAPS) given the specified commutation rules. Our experiments show that relaxed dependencies can indeed provide better solutions.

Non-Ancillary Mapping

The actual cost of ancillary qubits in practical quantum computing depends on the context. We provide optimal layout synthesis without any ancillary SWAPs as an option. Note that the resulting encoding may require more SWAP gates than when allowing ancillary qubits. This option can be encoded in our Two-Way SAT encoding, by simply restricting the SWAP gates to cases where both the physical qubits are already mapped.

3.2 Design Choices

Redundant Cardinality Constraints for Mapping Variables

Specifying Exactly-One constraints (EO) on logical qubits and At-Most-One constraints (AO) on physical qubits in the initial time step would be sufficient for correctness. Note that once the mapping variables are set in the initial time step, the information on bijectivity is propagated to next time steps, based on the chosen SWAP variables. However, observe that unrelated to the choice of SWAP variables, some invariants apply to mapping variables in all time steps. Essentially, the EO and AO constraints are orthogonal to the SWAP variable assignment. We observed that adding such redundant constraints at each time step significantly improved solving times. Apparently, this local information can be exploited by the SAT solver during clause learning or unit propagation. Since we added mostly binary clauses, we conjecture that the improved solving time is due to improved unit propagation.

Two-Way Encoding vs Explicit CNOT Constraints

In this paper, we encoded CNOT constraints using a Two-Way encoding instead of explicit CNOT constraints. We chose Two-Way encoding for two main reasons. First, we encode transitive closure for predecessors/successors of CNOT gates. Explicit constraints for CNOT dependencies result in two challenges:

- Specifying that the predecessors (successors) of a CNOT gate can not be scheduled in later (earlier) time steps results in a quadratic blow-up in the CNOT gates.
- Specifying that the predecessors (successors) of a CNOT gate must be scheduled in earlier (later) time steps results in long clauses.

On the other hand, the Two-Way encoding expresses this bidirectional propagation implicitly using clauses linear in the number of CNOT gates. Second, in incremental solving, in each iteration, we need to specify that the goal is reached in the final time step. Using the Two-Way encoding, we can simply specify that no CNOTs are delayed in the final step. If we encoded the CNOT constraints explicitly, we would need to specify that every CNOT is scheduled at some time step. To avoid large clauses, one would need to use auxiliary variables similar to advanced/delayed variables to keep track of the scheduled CNOTs across time steps. In our tool, we provide the encoding with explicit CNOT constraints as an option.

4 Experimental Evaluation

4.1 Experiment Design

We have extended our tool Q-Synth v1 (Quantum Synthesizer) to include the Two-Way Parallel SAT encoding. We provide an open-source tool Q-Synth2⁴ that implements the SAT encoding and the additional options. For any option chosen, our tool synthesizes

⁴ Q-Synth v2 tool with source code, benchmarks, and scripts <https://github.com/irfansha/Q-Synth>

a mapped circuit with the (provably) optimal number of additional SWAP+bridge gates. We use pysat [9] for generating and solving SAT instances incrementally. For cardinality constraints, we use the sequential counter from pysat. As a backend for our experiments, we use Cadical-1.53 [2], a state-of-the-art SAT solver. One can easily experiment with other SAT solvers in our tool using the pysat interface. When optimizing the SWAP count, our tool refutes all $k - 1$ SWAP+bridge additions if k is optimal. We report a timeout if an optimal solution is not found within the time limit. We check equivalence between the original circuits and our mapped circuits with QCEC⁵ [5] for correctness.

We design 3 experiments. Our goal is to investigate the effectiveness of our SAT encoding compared to the current leading tools. We also compare various additional techniques discussed in 3.1. For comparison, we consider state-of-the-art tools TB-OLSQ2 (near-optimal) [16] and Qiskit’s SABRE (heuristic). For TB-OLSQ2, we enable the best options i.e., SWAP optimization and upper bound computation by SABRE, with z3 (v4.12.1.0) [8] as the backend. TB-OLSQ2 can provide intermediate non-optimal results. We only report the final (near-optimal) solution when it terminates. If the tool does not terminate within the time limit, we report it as a timeout. For SABRE, we use the first 1000 seeds for the SABRE layout and take the minimum SWAPs generated by any seed. We also compare our results with other leading tools in Section 1.

Experiment 1: Standard Benchmarks on Large Platforms

We consider the standard benchmarks from papers [16, 26, 28] with 23 instances in total. The benchmark set contains circuits of up to 54 qubits and 270 CNOT gates. The circuits are mapped to the current NISQ processors, Sycamore with 54 qubits [1], Rigetti with 80 qubits⁶ and Eagle with 127 qubits [6]. We compare with the tools TB-OLSQ2 and SABRE, with a time limit of 12000 seconds (3hr 20 minutes) for each instance and an 8 GB memory limit.

Experiment 2: Deep VQE Benchmarks on Large Platforms

From our experiments and also consistent with the literature [16], most of the benchmarks from Experiment 1 need at most 10 SWAPs on standard platforms. To investigate the performance on deep circuits that need many SWAPs, we use a set of 10 random circuits composed using operators from the Variational Quantum Eigensolver (VQE) algorithm presented in [17]. Our second benchmark set consists of 10 (8 qubit) circuits with up to 79 CNOT gates. Due to many interactions between the qubits, the number of SWAPs needed to map onto the standard quantum platforms is high. We use the same time and memory limits as in Experiment 1. In both Experiments 1 and 2, we denote a timeout with TO. Here we focus on comparison with TB-OLSQ2 and report both SWAP count and circuit depth.

Experiment 3: Effectiveness of Additional Functionality

In this experiment, we compare 4 combinations of SWAPs (S), bridges (B), and relaxed dependencies (R): 1) S 2) S+B 3) S+R 4) S+B+R. From our two benchmark sets, we consider all the circuits with 14 or fewer qubits and map them onto the standard Melbourne platform of 14 qubits. We give a time limit of 600 seconds (or 10 minutes) and an 8 GB of memory.

⁵ Munich Quantum Toolkit QCEC <https://github.com/cda-tum/mqt-qcec>

⁶ Rigetti Computing <https://www.rigetti.com>

Of the 24 instances generated, we drop qft_8 which times out in all 4 combinations. For the rest of the 23 instances, we report SWAPs+bridges for each combination. Note that every additional SWAP or bridge adds exactly 3 extra CNOTs to the mapped circuit.

4.2 Results

■ **Table 2** Experiment 1: Number of SWAPs required for mapping circuits with QS2: Q-Synth2 (SWAP-optimal), TO2: TB-OLSQ2 (near optimal), and SB: SABRE (heuristic) tools on different platforms. Syc: Sycamore (54), Rig: Rigetti (80), Eagle (127) and label or(3/6) represents a circuit “or” with 3 qubits and 6 CNOT gates.

platform (qubits): Circuit(q/cx) / Tool	Sycamore (54)			Rigetti (80)			Eagle (127)		
	QS2	TO2	SB	QS2	TO2	SB	QS2	TO2	SB
or(3/6)	2	2	3	2	2	2	2	2	2
adder(4/10)	0	0	0	0	0	0	2	2	2
qaoa5(5/8)	0	0	1	0	0	0	0	0	1
4mod5-v1_22(5/11)	3	3	4	3	3	5	3	3	6
mod5mils_65(5/16)	6	6	7	6	6	7	6	6	8
4gt13_92(5/30)	10	10	15	10	10	15	13	TO	15
tof_4(7/22)	1	1	3	1	1	11	3	3	5
barenco_tof4(7/34)	5	5	18	6	6	17	8	8	17
qft_8(8/56)	9	TO	15	TO	TO	12	TO	TO	23
tof_5(9/30)	1	1	3	1	1	5	3	3	12
mod_mult55(9/40)	6	6	9	7	8	16	12	TO	20
barenco_tof5(9/50)	6	6	10	8	8	19	12	TO	20
vbe_adder3(10/50)	7	7	8	8	8	14	10	10	33
rc_adder6(14/71)	TO	8	16	8	8	35	TO	TO	51
ising_model10(16/90)	0	0	0	0	0	0	0	0	0
queko(16/15)	0	0	1	0	0	2	0	0	0
queko(16/29)	0	0	5	0	1	12	2	2	14
queko(16/44)	0	0	7	0	1	25	2	2	37
queko(16/58)	0	0	12	0	1	20	4	TO	41
queko(16/87)	0	1	10	0	1	30	4	TO	36
queko(16/101)	0	0	18	0	1	43	TO	TO	36
queko(54/54)	0	1	12	1	1	31	TO	TO	47
queko(54/270)	0	1	183	TO	TO	302	TO	TO	428
Total solved of 23	22	22	23	21	21	23	18	13	23


Experiment 1

Table 2 reports the number of SWAPs added. Both on Sycamore and Rigetti, TB-OLSQ2 mostly reports optimal SWAP count (while not proving optimality). There are 10 instances where it reports near-optimal solutions i.e., only 1 extra SWAP gate or times out. The difference is more significant on the larger 127-qubit Eagle platform. Q-Synth2 solved 5 more instances optimally where TB-OLSQ2 times out. Figure 7a provides the scatter plot of the time taken by Q-Synth2 and TB-OLSQ2. Except for two instances with rcadder6 (14 qubits) on Sycamore and Rigetti, we significantly outperform TB-OLSQ2 on all three platforms. In

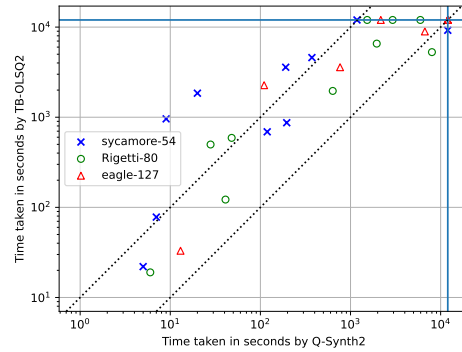
several instances, Q-Synth2 is one or two orders of magnitude faster while proving optimality. In the case of two instances with radder6, being a 14 qubit circuit, our tool takes time to refute the k-1 number of optimal SWAPs. While the heuristic tool SABRE always returns a mapping within 2 minutes, it also adds too many additional SWAPs. This observation is consistent with the literature [16].

■ **Table 3** Experiment 2: Additional SWAPs (s.) and Depth (d.) of mapped VQE circuits on different platforms with QS2: Q-Synth2 (SWAP-optimal) and TO2: TB-OLSQ2 (near optimal).

platform	Syc (54)				Rig (80)				Eagle (127)			
	QS2		TO2		QS2		TO2		QS2		TO2	
	s.	d.	s.	d.	s.	d.	s.	d.	s.	d.	s.	d.
Circuit(q/cx)												
vqe(8/18)	2	34	2	33	2	36	2	33	3	38	3	35
vqe(8/39)	4	65	5	62	6	68	7	63	7	65	9	64
vqe(8/40)	6	70	7	67	7	67	7	69	10	76	12	68
vqe(8/47)	8	85	10	83	10	84	10	83	14	92	17	86
vqe(8/48)	6	90	6	84	8	94	8	89	12	94	TO	TO
vqe(8/52)	9	90	11	87	11	90	TO	TO	TO	TO	TO	TO
vqe(8/63)	10	101	11	99	13	102	13	102	TO	TO	TO	TO
vqe(8/71)	TO	TO	16	111	15	112	18	113	TO	TO	TO	TO
vqe(8/78)	14	129	TO	TO	17	136	TO	TO	TO	TO	TO	TO
vqe(8/79)	11	149	11	146	16	151	TO	TO	TO	TO	TO	TO
Solved (10)	9		9		10		7		5		4	



(a) Experiment 1.



(b) Experiment 2.

■ **Figure 7** Scatter plots of time taken by TB-OLSQ2 and Q-Synth2.

Experiment 2

Table 3 reports the number of SWAPs added and circuit depth for the mapped VQE circuits. On all platforms, Q-Synth2 solves 15 more instances SWAP-optimally out of the 24 instances solved by either of the tools. TB-OLSQ2 in general reports better circuit depth compared to Q-Synth2. Interestingly in two instances, vqe(8/40) and vqe(8/71), Q-Synth2 reports better circuit depth. This shows that TB-OLSQ2 is near optimal in both SWAP additions and

circuit depth. Overall Q-Synth2 also reports near-optimal depth while optimizing additional SWAPs. Figure 7b shows the scatter plot of time for Experiment 2. Except for two instances with vqe(8/71) on Sycamore and Rigetti, Q-Synth2 significantly outperforms TB-OLSQ2.

■ **Table 4** Experiment 3: Number of SWAPs+bridges required for mapping deep VQE circuits on Melbourne platform (14-qubits) in 600 seconds with Q-Synth2 with combinations of S: Swaps, B: bridges, and R: relaxed dependencies.

Circuit(q/cx)	S	SB	SR	SBR	Circuit(q/cx)	S	SB	SR	SBR
or(3/6)	2	2	1	1	vqe(8/18)	2	2	2	2
adder(4/10)	0	0	0	0	vqe(8/39)	6	6	6	6
qaoa5(5/8)	0	0	0	0	vqe(8/40)	7	7	7	6
4mod5_22(5/11)	3	2	2	2	vqe(8/47)	8	8	8	8
mod5mils65(5/16)	6	4	4	4	vqe(8/48)	7	6	7	6
4gt13_92(5/30)	10	8	8	8	vqe(8/52)	10	10	10	10
tof_4(7/22)	1	1	1	1	vqe(8/63)	12	12	12	12
barencof4(7/34)	5	5	5	5	vqe(8/71)	13	12	13	12
tof_5(9/30)	1	1	1	1	vqe(8/78)	17	15	16	14
modmult55(9/40)	7	7	7	7	vqe(8/79)	15	13	14	13
barencof5(9/50)	6	6	6	6					
vbe_adder(10/50)	8	8	6	6					
rcadder6(14/71)	9	8	9	8					

Experiment 3

Table 4 reports the number of SWAPs+bridges on the Melbourne platform, with a time out of 10 minutes. Both bridges and relaxed dependencies can reduce the optimal SWAP+bridge additions. We observe that both techniques together can reduce CNOT count further. For instance, vqe(8/78) only needs 14 SWAPs+bridges i.e., 9 fewer CNOTs compared to only adding SWAPs. If we drop any of the options, the optimal CNOT count is higher.

4.3 Discussion

Comparison to OLSQ, OLSQ2 and TB-OLSQ2

In [16], authors showed that TB-OLSQ2 significantly outperforms both OLSQ and OLSQ2. Because of the lack of grouping in OLSQ and OLSQ2, the make-span of SMT instances generated is very large. Since our tool outperforms TB-OLSQ2, we do not report results from the other two directly. TB-OLSQ2 optimization routines can also be used in our tool to avoid hard unsatisfiable instances when optimality is not needed.

Comparison to Q-Synth v1 with Classical Planning

In [26], we showed that Q-Synth v1 based on Classical Planning outperformed both OLSQ and QMAP. While the approach scales well for mapping circuits up to 9 qubits onto 14 qubits platforms, larger circuits are out of reach. For instance, Q-Synth v1 timed out on rcadder6 (Table 4) with 3 hours. Q-Synth2 maps the same instance optimally within 5 minutes onto the 14-qubit Melbourne platform. Q-Synth v1 does not scale well to the other larger quantum platforms. Mapping individual CNOTs in Q-Synth v1 results in long plan length. As discussed in the same paper, long plan lengths increase the difficulty of planning.

Comparison to QMAP and SATMAP

QMAP [33, 31] employs an SMT encoding that grows exponentially with the number of physical qubits. Even using subarchitectures, QMAP is unable to map circuits greater than 7 qubits. SATMAP [19] on the other hand, encodes the Layout Synthesis as a MAXSAT problem to minimize the number of SWAPs. It allows the addition of one SWAP before every CNOT and uses MAXSAT solvers to minimize the number of SWAPs. As shown in [16], it produces suboptimal solutions and runs out of time even for moderate circuits.

Comparison with Dynamic Programming Approach

In [11], the authors provided an exact and a heuristic approach for adding SWAPs and bridge gates. With commutation rules, they showed that using bridges can further reduce the optimal CNOT additions. Our experiments are consistent with the authors' observations. Their exact approach already takes 12 minutes to map a 6-qubit circuit to 6-qubit platforms and grows exponentially with the number of qubits.

Comparison to SABRE

As observed in Experiment 1, it is clear that heuristic approaches such as SABRE add too many SWAPs. Adding many SWAPs not only increases the 2-qubit gate count but also increases circuit depth. However, heuristic approaches have their place in the quantum compilation pipeline. As the number of qubits on quantum processors increases, it is necessary to employ a hybrid strategy with heuristic and exact approaches. For instance, one could use SABRE to quickly find a reasonable initial mapping. Given such a mapping, one can synthesize a SWAP-optimal mapped circuit using Q-Synth2. As in TB-OLSQ2, we could use heuristic approaches to get quick upper bounds for near-optimal solving.

5 Conclusion

In this paper, we showed that parallel plans can be adapted to preserve SWAP optimality in layout synthesis. We have encoded the parallel planning problem directly in SAT. We propose a Two-Way encoding, in which information is propagated both forward and backward, for efficiency. We also demonstrated that our Two-Way SAT encoding is compatible with other techniques, like bridges and gate commutation rules.

The technique is implemented in the open-source tool Q-Synth2 for scalable and optimal layout synthesis of deep circuits. We can optimally map 8-qubit circuits that require up to 14 SWAPs onto a 127-qubit platform. We significantly outperform leading near-optimal tools while still guaranteeing that the resulting mapping is optimal.

References

- 1 Frank Arute et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019. doi:10.1038/s41586-019-1666-5.
- 2 Armin Biere, Katalin Fazekas, Mathias Fleury, and Maximilian Heisinger. CaDiCaL, Kissat, Paracooba, Plingeling and Treengeling entering the SAT Competition 2020. In *Proc. of SAT Competition 2020 – Solver and Benchmark Descriptions*, volume B-2020-1, pages 51–53. University of Helsinki, 2020. URL: <https://api.semanticscholar.org/CorpusID:220727106>.
- 3 Kyle E. C. Booth. Constraint programming models for depth-optimal qubit assignment and swap-based routing (short paper). In *29th International Conference on Principles and Practice of Constraint Programming, CP 2023, August 27-31, 2023, Toronto, Canada*, volume 280 of *LIPICs*, pages 43:1–43:10. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPICs.CP.2023.43.

- 4 Sebastian Brandhofer, Ilia Polian, and Kevin Krsulich. Optimal qubit reuse for near-term quantum computers. In *IEEE International Conference on Quantum Computing and Engineering, QCE 2023, Bellevue, WA, USA, September 17-22, 2023*, pages 859–869. IEEE, 2023. doi:10.1109/QCE57702.2023.00100.
- 5 Lukas Burgholzer and Robert Wille. Advanced equivalence checking for quantum circuits. *IEEE TCAD*, 40(9):1810–1824, 2021. doi:10.1109/tcad.2020.3032630.
- 6 Jerry Chow, Oliver Dial, and Jay Gambetta. Ibm quantum breaks the 100-qubit processor barrier. *IBM Research Blog*, 2, 2021. URL: <https://www.ibm.com/quantum/blog/127-qubit-quantum-processor-eagle>.
- 7 Andrew Cross, Ali Javadi-Abhari, Thomas Alexander, Lev Bishop, Colm A. Ryan, Steven Heidel, Niel de Beaudrap, John Smolin, Jay M. Gambetta, and Blake R. Johnson. Open quantum assembly language. *ACM Transactions on Quantum Computing Journal*, 2022. URL: <https://www.amazon.science/publications/open-quantum-assembly-language>.
- 8 Leonardo Mendonça de Moura and Nikolaj S. Bjørner. Z3: an efficient SMT solver. In *TACAS Proceedings*, LNCS 4963, pages 337–340. Springer, 2008. doi:10.1007/978-3-540-78800-3_24.
- 9 Alexey gnatiev, Antonio Morgado, and Joao Marques-Silva. PySAT: A Python toolkit for prototyping with SAT oracles. In *SAT*, pages 428–437, 2018. doi:10.1007/978-3-319-94144-8_26.
- 10 Stefan Hillmich, Alwin Zulehner, and Robert Wille. Exploiting quantum teleportation in quantum circuit mapping. In *ASPDAC '21*, pages 792–797. ACM, 2021. doi:10.1145/3394885.3431604.
- 11 Toshinari Itoko, Rudy Raymond, Takashi Imamichi, and Atsushi Matsuo. Optimization of quantum circuit mapping using gate transformation and commutation. *Integration*, 70:43–50, 2020. doi:10.1016/j.vlsi.2019.10.004.
- 12 Toshinari Itoko, Rudy Raymond, Takashi Imamichi, Atsushi Matsuo, and Andrew W. Cross. Quantum circuit compilers using gate commutation rules. In *ASPDAC*, pages 191–196. ACM, 2019. doi:10.1145/3287624.3287701.
- 13 Henry A. Kautz, David A. McAllester, and Bart Selman. Encoding plans in propositional logic. In *Proceedings of KR-96*, pages 374–384, November 1996. URL: <https://henrykautz.com/papers/plankr96.pdf>.
- 14 Henry A Kautz and Bart Selman. Planning as satisfiability. In *ECAI*, volume 92, pages 359–363, 1992. URL: <http://www.cs.cornell.edu/selman/papers/pdf/92.ecai.satplan.pdf>.
- 15 Gushu Li, Yufei Ding, and Yuan Xie. Tackling the qubit mapping problem for NISQ-era quantum devices. In *ASPLOS*, pages 1001–1014. ACM, 2019. doi:10.1145/3297858.3304023.
- 16 Wan-Hsuan Lin, Jason Kimko, Bochen Tan, Nikolaj Bjørner, and Jason Cong. Scalable optimal layout synthesis for NISQ quantum processors. In *DAC*, 2023. doi:10.1109/DAC56929.2023.10247760.
- 17 Marco Majland, Patrick Etenhuber, and Nikolaj Thomas Zinner. Fermionic adaptive sampling theory for variational quantum eigensolvers. *Phys. Rev. A*, 108:052422, November 2023. doi:10.1103/PhysRevA.108.052422.
- 18 Joao Marques-Silva, Ines Lynce, and Sharad Malik. Conflict-driven clause learning sat solvers. *Handbook of Satisfiability*, 336:133–182, 2021. doi:10.3233/FAIA200987.
- 19 Abtin Molavi, Amanda Xu, Martin Diges, Lauren Pick, Swamit Tannu, and Aws Albarghouthi. Qubit mapping and routing via MaxSAT. In *MICRO*, pages 1078–1091. IEEE, 2022. doi:10.1109/MICRO56248.2022.00077.
- 20 Yunseong Nam, Neil J. Ross, Yuan Su, Andrew M. Childs, and Dmitri Maslov. Automated optimization of large quantum circuits with continuous parameters. *npj Quantum Information*, 4(1), May 2018. doi:10.1038/s41534-018-0072-4.
- 21 Michael A. Nielsen and Isaac L. Chuang. *Quantum circuits*, pages 171–215. Cambridge University Press, 2010. doi:10.1017/CB09780511976667.008.
- 22 Tom Peham, Lukas Burgholzer, and Robert Wille. On optimal subarchitectures for quantum circuit mapping. *ACM Trans. on Quant. Computing*, 2023. doi:10.1145/3593594.

- 23 Jussi Rintanen, Keijo Heljanko, and Ilkka Niemelä. Planning as satisfiability: parallel plans and algorithms for plan search. *Artif. Intell.*, 170(12-13):1031–1080, 2006. doi:10.1016/J.ARTINT.2006.08.002.
- 24 Irfansha Shaik and Jaco van de Pol. Q-Synth. Software, version 2.0., IFD (Innovation Fund Denmark), swbId: swb:1:dir:be31c57364bd541c6b65afac80603e9004cf4008 (visited on 2024-07-31). URL: <https://github.com/irfansha/Q-Synth>.
- 25 Irfansha Shaik and Jaco van de Pol. Q-Synth v2.0 release. Software, version 2.0., IFD (Innovation Fund Denmark), swbId: swb:1:dir:be31c57364bd541c6b65afac80603e9004cf4008 (visited on 2024-07-31). URL: <https://github.com/irfansha/Q-Synth/releases/tag/Q-Synth-v2.0-SAT2024>.
- 26 Irfansha Shaik and Jaco van de Pol. Optimal layout synthesis for quantum circuits as classical planning. In *IEEE/ACM International Conference on Computer Aided Design, ICCAD 2023, San Francisco, CA, USA, October 28 - Nov. 2, 2023*, pages 1–9. IEEE, 2023. doi:10.1109/ICCAD57390.2023.10323924.
- 27 Amisha Srivastava, Chao Lu, Navnil Choudhury, Ayush Arunachalam, and Kanad Basu. Search space reduction for efficient quantum compilation. In *Proceedings of GLSVLSI-23*, pages 109–114. ACM, 2023. doi:10.1145/3583781.3590223.
- 28 Bochen Tan and Jason Cong. Optimal layout synthesis for quantum computing. In *IEEE/ACM ICCAD*, pages 137:1–137:9. IEEE, 2020. doi:10.1145/3400302.3415620.
- 29 Bochen Tan and Jason Cong. Optimality study of existing quantum computing layout synthesis tools. *IEEE Trans. Computers*, 70(9):1363–1373, 2021. doi:10.1109/TC.2020.3009140.
- 30 Davide Venturelli, Minh Do, Eleanor Rieffel, and Jeremy Frank. Temporal planning for compilation of quantum approximate optimization circuits. In *Proceedings of IJCAI-17*, pages 4440–4446, 2017. doi:10.24963/ijcai.2017/620.
- 31 Robert Wille, Lukas Burgholzer, and Alwin Zulehner. Mapping quantum circuits to IBM QX architectures using the minimal number of SWAP and H operations. In *DAC-19*, page 142. ACM, 2019. doi:10.1145/3316781.3317859.
- 32 Alwin Zulehner, Hartwig Bauer, and Robert Wille. Evaluating the flexibility of a* for mapping quantum circuits. In *Reversible Computation - 11th International Conference, RC 2019, Lausanne, Switzerland, June 24-25, 2019, Proceedings*, volume 11497 of *Lecture Notes in Computer Science*, pages 171–190. Springer, 2019. doi:10.1007/978-3-030-21500-2_11.
- 33 Alwin Zulehner, Alexandru Paler, and Robert Wille. An efficient methodology for mapping quantum circuits to the IBM QX architectures. *IEEE TCAD ICS*, 38(7):1226–1236, 2019. doi:10.1109/TCAD.2018.2846658.