Planning for Compilation of a Quantum Algorithm for Graph Coloring

Minh Do[1] and **Zhihui Wang**[2] and **Bryan O'Gorman**[3]
and **Davide Venturelli**[4] and **Eleanor Rieffel**[5] and **Jeremy Frank**[6]

Abstract. Recently, the problem of compiling general quantum algorithms for implementation on near-term quantum processors has been introduced to the AI community. Previous work demonstrated that temporal planning is an attractive approach for part of this compilation task, specifically, the routing of circuits that implement the Quantum Alternating Operator Ansatz (QAOA) applied to the MaxCut problem on a quantum processor architecture. In this paper, we extend the earlier work to route circuits that implement QAOA for Graph Coloring problems. QAOA for coloring requires execution of more, and more complex, operations on the chip, which makes routing a more challenging problem. We evaluate the approach on state-of-the-art hardware architectures from leading quantum computing companies. Additionally, we investigate applying the planning approach to qubit initialization as well as routing. Our empirical evaluation shows that temporal planning compares well to reasonable analytic upper bounds [20], and that solving qubit initialization with a classical planner generally helps temporal planners in finding shorter-makespan compilations for QAOA for Graph Coloring. These advances suggest that temporal planning can be an effective approach for more complex quantum computing algorithms and architectures.

1 Introduction

Quantum computers apply quantum operations, called quantum gates, to qubits, the basic memory unit of quantum processors. Quantum algorithms are often specified as quantum circuits on idealized hardware in which perfect gates can be applied to any set of qubits, whereas physical hardware has various constraints and imperfections. In practice, these idealized quantum circuits must be compiled to specific hardware. One common way of overcoming the restricted connectivity is by adding additional gates that route qubit states to locations where the desired gate can act on them. Compilations that minimize the overall duration return results more quickly, but more importantly are vital to obtain results on near-term quantum hardware that does not support significant quantum error correction: decoherence effects can destroy the computation in a short time.

Recently, the use of temporal planners to compile quantum circuits was explored for QAOA applied to the MaxCut problem [23]; machine operations were modeled as PDDL2.1 durative actions, enabling domain-independent temporal planners to find a parallel sequence of conflict-free operations to implement the high-level quantum algorithm. Several state-of-the-art temporal planners were used to show empirically that temporal planning is a promising approach to compile circuits of various sizes to a model hardware chip featuring the essential characteristics of newly emerging quantum hardware. Building upon this work, several subsequent works have utilized different techniques to more effectively solve the same routing instance set. In [4], the authors extended the planning-based approach by integrating it with a constraint-programming solver to further improve the plan quality; and [19, 21] explored greedy randomized search and genetic algorithms. While those approaches provide improved results, they require building domain-dependent heuristics or encodings that lack the flexibility of the model-based domain-independent temporal-planning approach introduced in [23], which can work on different classes of routing instances with a variety of hardware constraints and configurations.

In this paper, we expand the scope of the routing problem with a new target domain: QAOA for Graph Coloring [14, 25]; specifically, the optimization variant in which the number of properly colored edges is maximized. Compiling QAOA for Graph Coloring is different, harder, and more general than compilation of QAOA for MaxCut because of: (i) the existence of mix operations on two logical qubits (qstates); this leads to much more contention for resources on the gate-model hardware; and (ii) containing more elements of general compilation instances, and thus the ability to solve this instance class is key to future efforts to effectively utilize real-world gate-model quantum computers. Our main contributions over previous work are:

- *New instance class*: while our previous work concentrated on routing of QAOA for MaxCut, here we investigate routing of QAOA for Graph Coloring. Compiling Graph Coloring into physical circuit requires a different set of gates, which include 'hybrid' gates, and more complex ordering between them, making the compilation task much more complicated. To our knowledge, we present a compilation study on the most complex Noisy-Intermediate Scale Quantum (NISQ) optimization algorithm application to date.
- *More diverse physical hardware architectures*: while previous work have been using an earlier hypothetical model from Rigetti, the computational results for this paper were conducted using hardware graphs and gate durations that are closer to the real hardware that Google, IBM, and Rigetti are building.
- *Qubit initialization*: We present initial research results singling out the problem of qubit initialization (QI), which is a sub-problem related to routing. Our approach of using classical planning to solve QI improves the performance of all tested temporal planners across a variety of problem setups.

The paper is structured as follows. The next section provides background on quantum circuit routing. Then, in Section 3, we outline the

[1] KBR & NASA ARC, USA, minh.do@nasa.gov
[2] USRA & NASA ARC, USA, zhihui.wang@nasa.gov
[3] UC Berkeley & NASA ARC, USA, bogorman@berkeley.edu
[4] USRA & NASA ARC, USA, davide.venturelli@nasa.gov
[5] NASA ARC, USA, eleanor.rieffel@nasa.gov
[6] NASA ARC, USA, Jeremy.D.Frank@nasa.gov

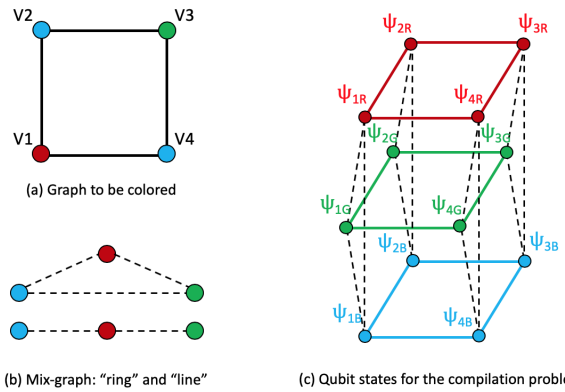

Figure 1: Routing for GraphColoring: coloring a square with 3 colors.

problem of circuit routing for QAOA on Graph Coloring and Section 4 describes how it can be modeled as a temporal planning problem using the PDDL2.1 standard modeling language. In Section 6, we describe different experiments and results showing the viability of our approach. Section 7 concludes the paper and outlines some of our future work directions.

2 Quantum Circuit Routing

General quantum algorithms historically have been described in an idealized architecture in which a gate can act on any subset of qubits. However, in an actual superconducting qubit device, such as the latest chips manufactured by IBM, Rigetti Computing, Google and Intel [7, 18, 2], physical constraints impose restrictions on the sets of qubits on which gates can be performed. Recently, a significant number of approaches have been explored for compiling idealized quantum circuits to realistic quantum hardware with a specific focus on "circuit routing"[7] (i.e., swap gate insertion strategies) in NISQ devices, targeting algorithms that could be run in the near-term [8, 26, 17, 21, 19, 13, 16].

For superconducting qubit architectures, qubits in these quantum processors can be thought of as nodes in a planar graph, with 2-qubit quantum gates associated with edges and 1-qubit quantum gates associated with nodes. Gates that operate on distinct sets of qubits may be able to operate concurrently, subject to additional restrictions, such as requiring the sets involved with concurrent gates to be non-adjacent. Furthermore, there are different types of quantum gate, each taking different duration that is dependent on the specific physical implementation. In order for the computation specified by the idealized circuit to be carried out, a particular type of 2-qubit gate, the *swap* gate, is often applied to exchange the state of two qubits. A sequence of swap gates moves the logical states of two distant qubits to a location where a desired gate can be applied. Swap gates may be available only on a subset of edges in the hardware graph, and swap duration may depend on where they are located. For the purposes of this study, we will consider the case in which swap gates are available between any two adjacent qubits on the chip and all swap gates have the same duration; the more general cases are a straightforward generalization.

3 Circuit Routing for Graph Coloring

In this paper, the Graph Coloring problem we will investigate is the *vertex coloring* problem in which all n vertices $v \in V$ of a given graph $\mathcal{G} = \{V, E\}$ are colored with k colors. The objective function is to maximize the number of edges $e = \langle v_1, v_2 \rangle \in E$ where vertices v_1 and v_2 are colored differently. This problem exemplifies the combinatorial structure of many scheduling and asset-allocation problems in industry and computer science research. Figure 1(a) shows a concrete example in which a square is properly colored with 3 colors.

[7] In previous work, we referred to this as "quantum circuit compilation" (QCC).

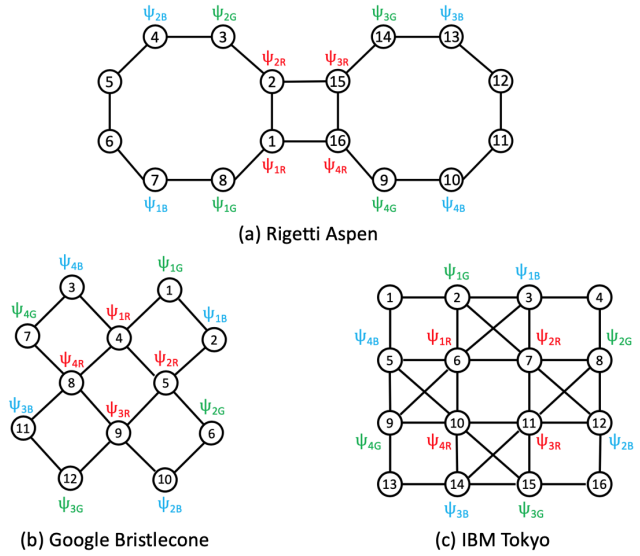

Figure 2: Example hardware chip layout from different companies.

The quantum algorithm that we follow is a variant of the "Quantum Alternating Operator Ansatz" [14], a generalization of the "Quantum Approximate Optimization Algorithm" (QAOA) [11], applied to the NP-hard Graph Coloring problem.

The Ideal Circuit: the idealized QAOA circuit has been studied in [14] and [25]. We refer the quantum-computing skilled reader to those references for understanding the algorithm, focusing here only on its compiler-level implementation. It is specified by two types of 2-qubit gate, the *phase separation* (PS) gate and the MIX gate, which need to be applied to a set of instance-specific problem *goals*[8]. Each goal specifies a pair of *qubit states* (qstate), the information content of a qubit, that must have a PS or MIX "goal" gate applied to them. For the Graph Coloring problem, each qstate represents a *vertex-color* combination. Thus, for our leading example problem of coloring a graph of 4 vertices ($V = \{V_1, V_2, V_3, V_4\}$) with 3 colors (red [R], green [G], and blue [B]), there are a total of $4 \times 3 = 12$ qstates: $\Psi = \{\psi_{1R}, \psi_{2R}, \ldots, \psi_{4B}\}$, illustrated in Figure 1(c). The idealized Graph Coloring circuit is specified as follows:

- *PS gate requirements*: for QAOA for GraphColoring, there should be one PS gate between any pair of qstates ψ_{iC} and ψ_{jC} that: (1) represent the same color C; (2) participate in an edge $e = \langle V_i, V_j \rangle \in E$. We denote the application of a PS gate as, e.g., PS(ψ_{1R}, ψ_{2R}), PS(ψ_{1R}, ψ_{4R}), PS(ψ_{3B}, ψ_{4B}).
- *MIX gate requirements*: one parameter of the QAOA for Graph-Coloring is the "mix-graph" $\mathcal{G}_{mix}$ with: (1) nodes representing different colors; and (2) edges representing which pairs of colors require MIX gates. In essence, mixing operators generate transitions between states representing different colors of the same vertex. Efficient mixing plays an important role in the quantum computation by achieving constructive interference that leads to

[8] Additional single qubit gates are also required in graph coloring but they are not relevant for routing since they can be executed without constraints at zero duration, so we will disregard them.

a good solution. A complete mixing-graph would allow all colors to transit to each other faster than a minimally-connected mixing graph (a chain) where a color only transits to its neighboring colors in one step [25]. On the other hand, a denser mixing-graph can require more SWAP gates, leading to a longer circuit execution time. Figure 1(b) shows two examples of a mix-graph for the three colors used in our example: (1) "ring": each color is connected to the other two. If we remove one edge (e.g., $Green - Blue$), then we have a (2) "line" mix-graph. Given a mix-graph $\mathcal{G}_{mix}$, the requirement is to achieve one MIX gate for each pair of qstates ψ_{iC_j} and ψ_{iC_k} that: (1) represent the same vertex V_i; (2) are associated with two different colors C_j and C_k such that $\langle C_j, C_k \rangle$ is an edge in $\mathcal{G}_{mix}$. We denote the application of a MIX gate as, e.g, MIX(ψ_{1R}, ψ_{1G}), MIX(ψ_{1R}, ψ_{1B}), ..., MIX(ψ_{4G}, ψ_{4B}).

- *Goal orderings*: constraints on the circuit structure of QAOA enforce orderings between the PS and MIX gates: all PS gates involving a certain qstate ψ_{iC_j} should be completed before any MIX gate involving ψ_{iC_j} can be executed. Let's take the qstate ψ_{1R} as an example: the two PS gates PS(ψ_{1R}, ψ_{2R}) and PS(ψ_{1R}, ψ_{4R}) need to be completed before either of the two MIX gates: MIX(ψ_{1R}, ψ_{1G}) or MIX(ψ_{1R}, ψ_{1B}) can start. Note that goal ordering constraints do not enforce that *all* PS gates need to be completed before the MIX gates can start.

Architecture-specific operations: as described in Section 2, while the idealized quantum circuit for GraphColoring contains the set of goal PS and MIX gates along with their orderings and the assumption that any goal gate can be applied anytime, they need to be "compiled" into architecture-specific operations that can actually be executed on a particular quantum chip that has numerous hardware constraints.

Figure 2 shows several examples: (a) a 16-qubit Aspen chip layout by Rigetti Computing; (b) a 12-qubit section from a 72-qubit Bristlecone chip by Google; and (c) a 16-qubit section of IBM's 20-qubit Tokyo architecture.

The set of operations for GraphColoring is significantly more complicated compared to the previous work on planning for routing of MaxCut [23, 4] where there are only three operations: 2-qubit PS operation, 1-qubit MIX operation, and 2-qubit SWAP operation. The increase in complexity arises also due to the opportunity of synthesis of 'multi-purpose' or 'hybrid' operations that accomplish multiple objectives (e.g., SWAP + MIX). More specifically, the following types of operations need to be considered by the compiler to tackle QAOA for GraphColoring on those architectures:

- *PS and MIX operations*: direct implementations of the ideal circuit PS and MIX gates that can be in principle applied to any pair of qstates residing on two qubits that are adjacent/interconnected on the hardware chip.
- *SWAP operation:* swaps the locations of two qstates residing on two interconnected qubits. For example, taking the layout of qstates on the Rigetti chip (Figure 2(a)): SWAP($\langle \psi_{2B}, q_4 \rangle, \langle \psi_{2G}, q_3 \rangle$) leads to: $\langle \psi_{2B}, q_3 \rangle$ and $\langle \psi_{2G}, q_4 \rangle$.
- *MOVE operation*: a variation of the SWAP operation that instead of swapping the locations of the two adjacent qstates, it moves a qstate to an adjacent empty qubit. For example (Figure 2(a)): MOVE($\langle \psi_{2B}, q_4 \rangle, q_5$) leads to $\langle \psi_{2B}, q_5 \rangle$, leaving q_4 empty[9].
- *SWAP-PS operation*: this operation combines the effects of the SWAP and PS operations. Thus, SWAP-PS($\langle \psi_{1R}, q_1 \rangle, \langle \psi_{2R}, q_2 \rangle$) in Figure 2(a) will switch the locations of ψ_{1R} and ψ_{2R} like the SWAP operation but also accomplish PS(ψ_{1R}, ψ_{2R}). Since the SWAP and SWAP-PS operations may have different durations depending on the particular physical connection, the planner needs to decide if it's beneficial to use one over the other at a particular location on the chip. For example, for a pair of qstates that do not have a PS goal gate requirement, the SWAP-PS gate can be applied, but its effect is just like conducting a SWAP gate between them. Thus: SWAP-PS($\langle \psi_{2B}, q_4 \rangle, \langle \psi_{2G}, q_3 \rangle$) has the same effect as SWAP($\langle \psi_{2B}, q_4 \rangle, \langle \psi_{2G}, q_3 \rangle$) since there is no goal gate requirement PS(ψ_{2B}, ψ_{2G}).
- *SWAP-MIX operation*: similar to the SWAP-PS operation, this one combines the effects of the SWAP and the MIX operations. However, unlike SWAP-PS that can be applied to any pair of qstates on adjacent qubits, the SWAP-MIX operation can only be applied between qstates representing nodes in the same mix-graph (Figure 1(b)). Let's assume the "line" mix-graph in Figure 1(b), which has two connections $B \leftrightarrow R$ and $R \leftrightarrow G$, then: (1) SWAP-MIX($\langle \psi_{1R}, q_1 \rangle, \langle \psi_{1G}, q_8 \rangle$) has the combined effects of SWAP and MIX gates; (2) SWAP-MIX($\langle \psi_{1B}, q_7 \rangle, \langle \psi_{1G}, q_8 \rangle$) has the same effect as SWAP gate (but can have a shorter makespan to make it a better option than SWAP) since ψ_{1B} and ψ_{1G} are not connected on, but belong to the same, the mixgraph; (3) SWAP-MIX($\langle \psi_{1R}, q_1 \rangle, \langle \psi_{2R}, q_R \rangle$) are not allowed since they are not connected on a mix-graph.

Problem definition: Given an idealized circuit, comprised of PS and MIX gates, used to define a QAOA quantum algorithm for GraphColoring, the circuit routing problem is to find a new architecture-specific circuit that implements the idealized quantum circuit, utilizing the architecture-specific PS, MIX, SWAP, MOVE, SWAP-PS, and SWAP-MIX operations as required. The objective is to minimize the overall duration to execute all operations in a new circuit.

4 Model Circuit Routing for GraphColoring as a Temporal Planning Problem

Planning is the problem of finding a conflict-free set of actions and their respective execution times that connects the *initial-state* I and the desired *goal state* G. We now introduce some key background concepts for the routing-as-temporal planning problem.

Planners: A planner is a software that takes as input a specification of domain and problem descriptions and returns a valid plan if one exists. At the abstract level, the planner needs to solve the QAOA compilation problem described in previous section: taking as input the required PS and MIX gates and utilizing the architecture and problem-specific operations to build a plan achieving all those gates.

Planning Domain Description Language (PDDL): is the de facto standard modeling languages used by many domain-independent planners. We use PDDL 2.1, which allows the modeling of temporal planning formulations in which every action a has duration d_a, starting time s_a, and end time $e_a = s_a + d_a$. Action conditions are required to be satisfied either (1) instantaneously at s_a or e_a or (2) to be true starting at s_a and remain true until e_a. Action effects may instantaneously occur at either s_a or e_a. Actions can execute when their temporally-constrained conditions are satisfied; and when executed will cause state-change effects. The most common objective function in temporal planning is to minimize the plan *makespan*, i.e., the shortest total plan execution time. This objective matches well with the objective of our targeted routing problem: minimizing the total circuit execution time (i.e., circuit depth).

Modeling Routing for GraphColoring in PDDL 2.1: Following the software structure of the end-to-end compiler tool-chain presented in [22], at the highest level, we need to take as input:

[9] This operation does correspond to a SWAP operation in the real chip, where empty qubits don't exist. It is defined only for modeling convenience.

```
(:durative-action swap_mix_at_q1_q2
    :parameters (?s1 - qstate ?s2 - qstate)
    :duration (= ?duration 1.0)
    :condition (and (at start (located_at_q1 ?s1))
                    (at start (located_at_q2 ?s2))
                    (at start (PS_phase_completed ?s1))
                    (at start (PS_phase_completed ?s2))
                    (at start (mixgraph_edge ?s1 ?s2))
                    (at start (not (mixed ?s1 ?s2))))
    :effect (and (at start (not (located_at_q1 ?s1)))
                 (at start (not (located_at_q2 ?s2)))
                 (at end (located_at_q1 ?s2))
                 (at end (located_at_q2 ?s1))
                 (at end (mixed ?s1 ?s2))
                 (at end (mixed ?s2 ?s1))))
```

Figure 3: PDDL model of the SWAP-MIX operation.

- Qstates and their relations: this in turn encapsulates the graph to be colored $\mathcal{G}$ (Figure 1(a)) and the "mix-graph" $\mathcal{G}_{mix}$ (Figure 1(b)).
- The $\mathcal{G}_{machine}$ graph representing the hardware chip layout (Figure 2).

and turn them into *objects*, *predicates*, *actions*, *initial state*, and *goal state* of a temporal planning problem such that a valid plan represents a parallel sequence of architecture-specific operations (see Section 3) enabling all required PS and MIX gates.

Objects: like previous approaches in planning for routing, we model qstates as PDDL objects. Physical qubits and the connections in $\mathcal{G}_{machine}$, $\mathcal{G}_{mix}$, and $\mathcal{G}$ graphs are modeled implicitly within the list of predicates and action descriptions.

Predicates: the following facts are represented by PDDL predicates:

- *located_at*: if a given qstate is located at a given qubit.
- *empty*: if a given qubit is empty. This predicate enables the MOVE gate (see Section 3).
- *ps* and *mix*: if PS or MIX gate has been accomplished between a pair of qstates.
- *edge*: if two qstates represent an edge in the graph G to be colored. This predicate serves as the precondition of the PS gate.
- *mixgraph_edge*: if a given pair of qstates make up an edge in a mix-graph. This predicate serves as a precondition of the MIX gate and the SWAP-MIX gate to enable "mix" as action effect.
- *same_mixgraph*: if a given pair of qstates belong to the same mix-graph, but do not make up an edge in a mix-graph. This enables the SWAP-MIX action between those qstates, but does not lead to achieving "mix" as action effect.
- *PS_phase_completed*: if the PS phase for a given qstate is finished.

Actions: there are 6 action templates representing 6 architecture-specific operations described in Section 3: PS, MIX, SWAP, MOVE, SWAP-PS, and SWAP-MIX. These actions act on the edges of the graph $\mathcal{G}_{machine}$ with appropriate qstates as action parameters. As outlined in Section 3, the SWAP-MIX action's "mix" effects condition on whether or not the two involved qstates make up an edge in $\mathcal{G}_{mix}$. Given that many temporal planners can not handle conditional-effect, to allow us to use a wider range of planners, we compile it into two deterministic actions: (1) SWAP-MIX: operating on two qstates making up an edge in a mix-graph (i.e., having *mixgraph_edge* as a condition) and combine both MIX and SWAP action effects; and (2) SWAP-MIX-AS-SWAP: operating on two qstates belong to the same mix-graph but there is no edge connecting them (i.e., having *same_mixgraph* and $\neg mixgraph\_edge$ as its conditions). We also introduce an auxiliary instantaneous action DonePS(ψ) to specify that a qstate ψ is done with the PS phase and is ready to move on to the MIX phase. In our example shown in Figure 1, DonePS(ψ_{1R}) can be executed when PS(ψ_{1R}, ψ_{2R}) and PS(ψ_{1R}, ψ_{4R}) are completed and it in turn enables any MIX or SWAP-MIX operation involving ψ_{1R}, e.g. MIX(ψ_{1R}, ψ_{1B}) through its $PS\_phase\_completed$ effect (see Figure 3).

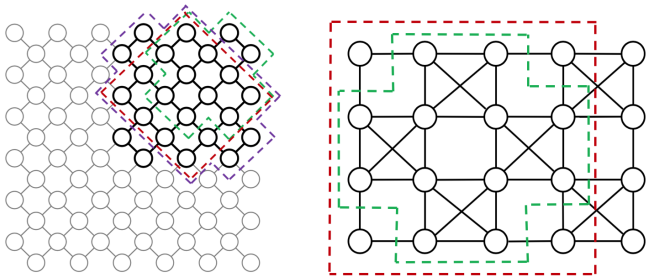

(a) Google 72-qubit Bristlecone chip (b) IBM 20-qubit Tokyo chip

Figure 4: Google and IBM full chip configurations and the sections of 12 (green), 16 (red), 20 (purple), and 24 qubits used for our empirical evaluation.

Initial state: the initial state declares the initial values of all predicates: *located_at* and *empty* specify the initial locations of all qstates (and if some physical qubits are empty); *mixgraph_edge* and *same_mixgraph* values capture $\mathcal{G}_{mix}$ configuration; and *edge* values represent the input graph $\mathcal{G}$ to be colored.

Goal state: the goal state specify the list of MIX gates that need to be achieved. The *donePS* conditions of the MIX actions (see Figure 3) capture the ordering constraints between PS and MIX gates and ensure that all requisite PS gates are also achieved.

5 Qubit Initialization

Qubit initialization (QI) is the problem of assigning the initial locations of all qstates on the chip. Figure 2 shows examples of QI on the three machine configurations. Two approaches for QI were explored in the previous routing as Temporal Planning for MaxCut work [4]: (1) random initialization of qubit; and (2) for each qstate q, add an action INIT(q, p) to locate qstate q on the physical qubit p; all INIT actions need to be carried out before any other action can start.

Basically, the second approach combines QI with routing, resulting in a Routing-I problem [4], with the hope that current state-of-the-art temporal planners would be able to find good initial locations for all qstates that support finding lower makespan plan. While previous work compared different temporal planners on solving Routing-I for MaxCut, they didn't report the comparison between random initialization and Routing-I. Our current investigation for Routing as Temporal Planning for Graph Coloring shows that using a temporal planner to solve the Routing-I problems is not a clear-cut better option than random initialization (see results in Section 6.3), and it doesn't perform well compared to careful manual initialization utilizing domain knowledge about the problem structure and the hardware chip layout. We believe this is due to: (1) the large number of possible initial configuration (for a n-qubit chip, the number of possible initial states is $n!$, modulo potential symmetries); and (2) since initialization needs to be fully done before gate composition begins, it is difficult for the existing planners to give a good heuristic estimate on which initial state is a good one to start from since they are furthest states from the goals. The farther a given state from the goals, the harder to get a good heuristic estimate for that state.

Qubit initialization as classical planning: instead of solving the combined Routing-I problem, we can heuristically solve the QI problem separately. Thus, we try to find the initial locations I of qstates so that

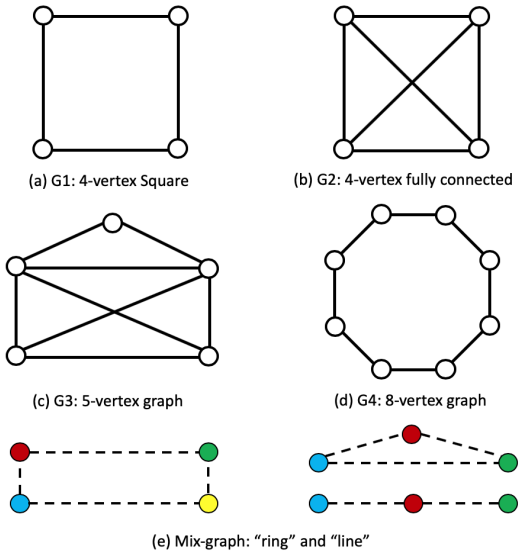

Figure 5: The benchmark instance set: Graphs to be colored (a-d) and mix-graph configurations (e).

the subsequent routing problem will yield shorter makespan plans. As described in Section 3, the routing for Graph Coloring require each involved qstate to conduct all related PS gates, and then all related MIX gates. Since the PS gates will be done prior to the MIX gate, one heuristic is to initialize qstates so that if a pair of qstates ψ_1, ψ_2 requires PS(ψ_1, ψ_2) as a goal, then they would be initially allocated on two qubits that are close to each other on the hardware chip. In our leading example, shown in Figure 1 and 2, we would like to put the two qstates ψ_{1R} and ψ_{2R} close initially to minimize the number of SWAP gates needed before the required PS(ψ_{1R}, ψ_{2R}) gate can be conducted.

Let $dis(q_i, q_j)$ be the shortest duration to move a given qstate ψ_i located on q_i and qstate ψ_j located on q_j to two qubits that are connected so that PS(ψ_i, ψ_j) can be conducted (can be efficiently precomputed using a variation of the all-pair-shortest-path algorithm); and E_{PS} is the set of state pairs $\langle \psi_i, \psi_j \rangle$ such that PS(ψ_i, ψ_j) is a goal requirement. We then can map the QI problem into a quadratic assignment of finding the mapping function $f : V \rightarrow Q$ that minimize the total distances of all pairs of qstates that are involved with a PS gate:

$$minimize \sum_{\langle \psi_i, \psi_j \rangle \in E_{PS}} dis(f(\psi_i), f(\psi_j))$$

The quadratic assignment problem [9] is NP-hard. Therefore, instead of trying to find an optimal solution, we will heuristically solve it by first setting up a new planning problem P' as follows:

1. Remove all temporal aspects of the original planning problem: turn it into a simpler classical planning problem.
2. Remove all goals related to the MIX gate requirements.
3. Remove all actions except SWAP, PS, and SWAP-PS.
4. Set the non-temporal SWAP and SWAP-PS action cost in P' to be equal to the duration of the temporal SWAP and SWAP-PS actions in P and the PS action cost to be 1.

We then use a classical planner to find a solution π for P' with the objective function of minimizing the total plan cost. We then use the final locations of all qstates after executing π as the initial qstate allocation for the original temporal planning problem P. The hypothesis is that since the goal of P' is to achieve all PS gates, qstates that are engaged in the PS gates in P would likely be close to each other when π is done executing; their final locations are thus good initial locations for the original planning problem P where the same set of PS gates need to be accomplished before the subsequent MIX gates.

Given the simplified problem and the more abundant of classical planners, compared to temporal planners, the qubit-initialization-as-planning problem should be significantly easier to solve compared to the original Routing as Temporal Planning problem.

Table 1: Plan quality comparison in solving problems shown in Figure 5 with different hardware architectures shown in Figure 2 and 4. Qubit initialization is done manually using expert knowledge, similar to those in Figure 2. **bold** values indicate the best overall makespan. "-" indicates either the planner ran out of time before finding a solution (OPTIC) or crashed (SGPlan).

	Graph	TFD	OPTIC	LPG	SGPLAN
Rigetti-12	G1R3	**28**	**28**	31	61
	G1L3	**27**	42	33	44
Google-12	G1R3	**22**	45	46	83
	G1L3	**21**	38	41	68
IBM-12	G1R3	**23**	37	31	51
	G1L3	**17**	30	36	39
Rigetti-16	G1R4	**31**	-	80	94
	G2R4	**74**	-	119	156
Google-16	G1R4	**19**	46	20	34
	G2R4	76	49	**37**	48
IBM-16	G1R4	43	-	26	**20**
	G2R4	79	58	49	**29**
Google-20	G3R4	**64**	-	86	106
IBM-20	G3R4	113	-	**94**	-
Google-24	G4R3	125	**64**	83	-

6 Empirical Evaluation

In this section, we will present preliminary empirical evaluation of different temporal planners applied to routing for Graph Coloring.

Hardware Architecture: We use the recently published hardware chip architecture layouts from Rigetti (Aspen), Google (Bristlecone), and IBM (Tokyo). The full-size chip from Rigetti is shown in Figure 2(a), and sections of the full chip from Google and IBM are shown in Figure 2(b) and (c) respectively. Figure 4 shows the entirety of the 72-qubit Bristlecone and the 20-qubit Tokyo architectures. In our experiment, we use the following specs: (1) Rigetti: 12-qubit section and the full 16-qubit chip (Figure 2(a)); (2) Google: 12-, 16-, 20- and 24-qubit sections of the Bristlecone 72-qubit architecture (Figure 4(a)); (3) IBM: 12- and 16-qubit sections and the full 20-qubit Tokyo chip (Figure 4(b)). Overall, the hardware architectures under consideration have 12–24 qubits.

Software: Previous work on routing for MaxCut used TFD [10], LPG [12][10], CPT [24], POPF [6], and SGPlan [5]. In this paper, we: (1) exclude CPT due to poor scalability (for both Graph Coloring and Max Cut); (2) replace POPF with the more modern code of another planner, OPTIC [3], in the same family. For solving the QI-as-planning problem, we use the classical planner FastDownward [15].

Problem specifications: We test our approach on a range of Graph Coloring instances with 4, 5, and 8 vertices and either the "line" or "ring" mix-graph. They are shown in Figure 5(a–d) as graphs G1–4. For example, graph G2R4 (Table 1, 2, and 4) means solving graph G2 in Figure 5 with the "Ring" mix-graph with 4 colors. The number of vertices n in the graph and the number of colors k in the mix-graph

[10] Since LPG is a stochastic planner, for each problem, we ran LPG 10 times and report its best result.

Table 2: Comparing against analytical bounds for special hardware architectures: grid and line. Highlight in **bold** are makespan values that are better than the analytical bound. Underlined values are the best makespan produced by any planner but is still higher than the analytical bound.

	Analyt. Bound	TFD	OPTIC	LPG	SGPLAN
4 × 3 Grid	19	20	35	**16**	**13**
4 × 4 Grid	20	30	56	**20**	38
5 × 4 Grid	24	54	116	79	46
8 × 3 Grid	35	**25**	53	36	-
4 × 3 Line	64	**51**	77	**48**	90
4 × 4 Line	80	**71**	116	107	177
5 × 4 Line	100	118	-	140	194
8 × 3 Line	128	**81**	-	157	267

will require the number of physical qubits $n \cdot k$ within the range of 12-24 qubits, as described in the previous paragraph.

Gate durations: For all hardware architectures and for all edges in the hardware chips, the following gate durations are used: dur(SWAP) = 4, dur(MOVE) = 4, dur(PS) = 3, dur(SWAP-PS) = 4, dur(MIX) = 1, and dur(SWAP-MIX) = 1. These are effective durations based on logical gate synthesis using common native gate sets, e.g., that available on Rigetti's chips [1].

Results were collected on a RedHat Linux VM running on a Macbook Pro with 4GB of RAM. The runtime limit was set to 600 seconds for problems involving 12–16 qubits and 1200 seconds for problems involving 20–24 qubits.

6.1 Solving Circuit Routing with Temporal Planner

Table 1 compares results between 4 temporal planners. Overall, TFD and LPG are the only two planners that can solve all problems within the time limit. In terms of solution quality, TFD is most often the best planner, especially for smaller problems. However, each planner has at least one case where it performs best. SGPlan generally returns the worst quality plan but it also performs the best in the two problems on IBM 16-qubit architecture. With regard to *ring* vs *line* mix-graph, as expected, the ring version of the 12-qubit problems normally has slightly longer makespan value than the line counterpart, given that it requires additional MIX goal gates. Between different hardware architectures, as expected the Rigetti machine has much fewer connections for the same chip size (12 or 16-qubit), which leads to fewer parallel routes to move qstates. This led to longer makespan plans, in general, compared to solving the same problems on the Google and IBM architectures. Comparing the Google and IBM architectures of the same size, we would expect planners should be able to find equal or shorter makespan plans given that they have the same overall shape except that the IBM ones have some additional connections. While that generally holds true for the 12-qubit version, looking at the results for the 16-qubit version we see TFD, OPTIC, and LPG perform worse on the IBM architecture than the Google architecture. It seems that the additional connections enlarge the planning search space and confuse the planners, thus the planners can't exploit the connectivity to find better quality plans. SGPlan is the only planner that show marked better performance on the IBM 16-qubit architecture.

6.2 Planner vs Analytical Bounds

Without another systematic approach to solve Routing for Graph Coloring to compare with, one question remains: *how good is the quality of solutions returned by temporal planners in this domain?*

Table 3: Improvement in makespan when solving a problem using qubit initialization provided by a classical planner (Section 5). Example: entry 8.1%(8/10) for cell IBM-16, G2R4, OPTIC means: among 10 randomly generated QIs for solving G2R4 on IBM-16, then OPTIC can solve 8 with random QI while can solve all 10 with QIs produced by the FastDoward classical planner solving the same random problems. Accross the 8 problems that are solved with both QI setups, the average makespan improvement is 8.1%. When the number of solved problem is not listed, it means all 10 are solved with both QI options. RED indicates entries where the FastDownward's QI performs worse than random QI.

	Graph	TFD	OPTIC	SGPlan
Rigetti-12	G1R3	17.9%	14.4%	19.1%
	G1L3	10.5%	2.7%	3.52%
Google-12	G1R3	9.7%	7.9%	-5.7%
	G1L3	16.3%	12.3%	13.1%
IBM-12	G1R3	9.2%	2.1%	10.9%
	G1L3	10.2%	18.4%	17.4%
Rigetti-16	G1R4	13.6%	23.8%(7/10)	15.7%
	G2R4	20.4%	24.6%(8/10)	18.1%
Google-16	G1R4	7.4%	5.3%	12.3%
	G2R4	17.5% (10/9)	7.9%(9/10)	11.3%
IBM-16	G1R4	11.2%	-3.9%	14.0%
	G2R4	6.4% (7/6)	8.1%(8/10)	7.2%
Google-20	G3R4	23.0%	2.6%(4/9)	23.7%
IBM-20	G3R4	12.6%	(2/3)	-
Google-24	G4R3	6.8% (9/9)	(0/2)	-

For comparison, we can use simple manually constructed solutions, which are currently available for certain problem setups, as an upper bound [20]. Recall that for coloring a graph of n vertices with k colors on a hardware chip, we use $n \cdot k$ qubits. For hardware that supports the gates described in Section 3, we can get valid plans for the two hardware architectures: an $n \cdot k$ "grid" layout and a $n \cdot k$-length "line" layout with the following makespans:

$$\text{makespan}_{LINE} \leq n \cdot \tau_{\text{SWAP\_PS}} + nk \cdot \tau_{\text{SWAP}} + k \cdot \tau_{\text{SWAP\_MIX}}$$
$$\text{makespan}_{GRID} \leq n \cdot \tau_{\text{SWAP\_PS}} + k \cdot \tau_{\text{SWAP\_MIX}}$$

This can be accomplished through a predetermined qubit initialization and operation sequence such that while there will be SWAP gates required for the "line" layout, the "grid" layout only needs combined gates (i.e., SWAP-MIX and SWAP-PS) to accomplish all goal gates. Table 2 shows that while each planner's performance is quite inconsistent, the best plan returned by the planning approach compare well with reasonable upper bounds for those "grid" and "line" chip layout.

6.3 Qubit Initialization as Classical Planning

To measure the effect of qubit initialization for Graph Coloring by solving a classical planning problem, for each of the 15 problem configurations described in the previous Section 6.1 and shown in Table 1, we: (1) generate 10 random qubit initializations; (2) use a classical planner FastDownward to solve the qubit initialization problem as described in Section 5; (3) solve two versions of the problem with: (i) random initialization and (ii) with initialization given by the solution returned by FastDownward. Table 3 shows the makespan improvement of using the same planner[11], solving (ii) over solving (i); averaged over 10 random problem for each configuration. We exclude LPG from this experiment due to its randomness nature. LPG

[11] We use the same running time limit as in collecting results for Table 1

Table 4: Comparing different initialization strategies. **M**: **M**anual initialization (results from Table 1); **I**: Solving the combined QI + Routing (Routing-I) problem in one temporal-planning run (see Section 5); **Random** and FastDownward(**FD**) are initialization setups explained in Table 3 with "AVG" and "Best" represent the average and best values across 10 random QIs. Values in RED show the best value across all setups; and in **bold** are values that are best for a given planner. Random vs FD qubit initialization: Light Yellow background indicates better makespan by FD while Light Cyan indicates better makespan for random QI.

		TFD						OPTIC						SGPlan						LPG
Problem		M	I	Rand. (AVG)	Rand. (Best)	FD (AVG)	FD (Best)	M	I	Rand. (AVG)	Rand. (Best)	FD (AVG)	FD (Best)	M	I	Rand. (AVG)	Rand. (Best)	FD (AVG)	FD (Best)	M
Rigetti-12	G1R3	28	-	61.2	53	50.1	42	28	76	73.7	63	62.7	44	61	89	94	69	75.9	**53**	31
	G1L3	27	-	52.3	44	46.3	39	**42**	62	65	49	62.9	48	**44**	84	73.4	62	70.5	59	33
Google-12	G1R3	22	47	39.8	33	35.8	27	45	40	48.7	**31**	43.7	34	83	77	60.9	**43**	65.4	52	46
	G1L3	21	50	38.4	34	31.9	25	38	28	43	34	37.7	**27**	68	69	62.1	46	52	**38**	41
IBM-12	G1R3	23	35	33.8	27	30.5	26	**37**	61	47.5	40	45.5	**37**	51	63	56.5	47	49.8	**43**	31
	G1L3	17	25	29.9	19	26.4	21	30	39	44.2	31	36.1	**25**	39	50	52.9	42	43.2	**38**	36
Rigetti-16	G1R4	31	77	62.4	49	51.7	44	-	66	78.4	52	58.3	**43**	94	118	119.3	92	97.4	**68**	80
	G2R4	74	-	83.2	65	64.5	56	-	85	94.5	76	72.1	**63**	156	140	139	105	111	**89**	119
Google-16	G1R4	19	-	42.4	36	39.2	31	46	51	58.9	47	55.8	**39**	**34**	94	86.3	57	73.5	55	20
	G2R4	76	-	98.3	55	79.8	**43**	49	68	71.8	61	65.7	**48**	**48**	103	112.3	83	98.1	75	37
IBM-16	G1R4	43	-	60.9	47	53.7	**31**	-	57	50.8	34	54.1	**32**	20	-	77.3	63	65.9	53	26
	G2R4	79	-	85	74	79.5	**70**	58	56	73.5	57	64.6	**51**	29	-	89.2	67	82.7	52	49
Google-20	G3R4	64	-	128.9	64	90.7	58	-	-	86	76	82.9	**62**	106	-	152.3	115	115.1	**88**	86
IBM-20	G4R4	113	-	111.7	82	96.3	48	-	-	77.5	**70**	81	76	-	-	-	-	-	-	94
Google-24	G4R3	125	-	155.7	134	143.6	**113**	64	37	-	-	72.5	61	-	-	-	-	-	-	83

would employ different random seeds when solving (i) and (ii), making the comparison not justifiable.

In essence, Table 3 shows, given a (random) qubit initialization, how different planners benefit from running a classical planner such as FastDownward as a pre-processing phase to potentially find better initial locations for all qstates. The results show that all three planners benefits from this step for a majority of testing scenarios: solving higher overall number of instances while experiencing makespan improvement with the qubit initialization calculated by FastDownward. There are only two case in which average makespan increases: OPTIC/ IBM-16/G1R4 and SGPlan/Google-12/G1R3. The reason is likely that the majority of the random initializations happen to be very good starting points and FastDownward plans move qstates to overall worse initializations. Note that the classical planning setup for QI approximates and relaxes the problem by: (1) ignore the MIX goals; (2) remove the temporal aspects (and thus finding a sequential non-parallel plan). Therefore, it's possible that final qstate locations provided by the FastDownward solution is worse than the starting random initialization.

Overall Comparison: Table 4 shows the makespan comparison between different settings for all planners. Overall, TFD consistently provide the best plan quality with manual initialization for smaller problems. For larger problems, where finding a good manual initialization is likely harder, different planners returns the best quality solutions for different problem settings. Taking each individual planner, the data show, as expected, the best results are from either manual initialization or provided by using the FastDownward (FD) planner. Between those two settings, while FD provides better results for TFD and OPTIC on larger problems, for SGT manual initialization provides better results for larger ones. While Routing-I, where temporal planner solves both the QI and routing problem in one planning model, would be the most convenient setup, the results show that existing temporal planners are still not being able to solve that problem effectively. Only one best makespan is found using this problem setup, provided by the OPTIC planner. It's also worth pointing out that while TFD is the most promising planner for routing, it performs worst on Routing-I with only 5/15 problems solved. Consistent with the results shown in Table 3, using QI result provided by the FastDownward planner, all three planners produce smaller makespan values (both average and best) compared to the random QI counterpart.

In summary, at large scale and for general instances, it is infeasible to manually solve the QI problem, and so we show that the process can be automated using a classical planner. The most promising planner is TFD and it should be the first one to try out. However, given that SGPlan is by far the fastest among all 4 tested planners, solving all tested problems within a few seconds, it's also worth running it besides other planners. While SGPlan is very inconsistent, it sometimes (see Table 4 and 2) produces plans with makespan values much lower than other planners.

7 Conclusion and Future Work

In this paper, we describe our recent investigation into using model-based planning technology to solve the quantum circuit routing for QAOA applied to the Graph Coloring problem: temporal planner for routing and classical planner for the QI problem. Our empirical evaluation shows that temporal planners can solve problems with diverse setup effectively, compare reasonably to the best known analytical bounds on special cases. Qubit initialization as classical planning, utilising the FastDownward planner, provides makespan improvement in majority of problem configurations and provide an attractive alternative to manual qubit initialization using domain expert knowlege.

There are several directions we are pursuing in future work. First, we are working on running some plans produced by the temporal planners described in this paper on actual, physical hardware chips. Second, as an alternative to utilizing classical planning, we are developing several other approaches to QI that do not use classical planning but other combinatorial optimization techniques such as MILP. Third, we would like to analyze better the synergy between different temporal planning algorithm and the problem structures that are most suitable to them. Lastly, several portfolio approaches have been showing good performances at the recent International Planning Competition (IPC); we would like to investigate their performance on the Routing for Graph Coloring domain, and compare them to our current set of temporal planners used in this paper.

REFERENCES

[1] Deanna M. Abrams, Nicolas Didier, Blake R. Johnson, Marcus P. da Silva, and Colm A. Ryan, 'Implementation of the xy interaction family by calibration of a single pulse', *to be published*, (2019).

[2] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, et al., 'Quantum supremacy using a programmable superconducting processor', *Nature*, **574**(7779), 505–510, (2019).

[3] J. Benton, Amanda Coles, and Andrew Coles, 'Temporal planning with preferences and time-dependent continous costs', in *Proceedings of the Twentieth-Second International Conference on Automated Planning and Scheduling (ICAPS-12)*, (2012).

[4] Kyle E. C. Booth, Minh Do, J. Christopher Beck, Eleanor Rieffel, Davide Venturelli, and Jeremy Frank, 'Comparing and integrating constraint programming and temporal planning for quantum circuit compilation', in *Proceedings of the Twenty-Eighth International Conference on Automated Planning and Scheduling (ICAPS2018)*, pp. 366–374, (2018).

[5] Y. Chen and B. Wah, 'Temporal planning using subgoal partitioning and resolution in sg-plan', *Journal of Artificial Intelligence Research*, **26**, 323 – 369, (2006).

[6] A. J. Coles, A. I. Coles, M. Fox, and D. Long, 'Forward-chaining partial-order planning', in *Proceedings of the Twentieth International Conference on Automated Planning and Scheduling (ICAPS-10)*, (2010).

[7] AD Corcoles, A Kandala, A Javadi-Abhari, DT McClure, AW Cross, K Temme, PD Nation, M Steffen, and JM Gambetta, 'Challenges and opportunities of near-term quantum computing systems', *arXiv preprint arXiv:1910.02894*, (2019).

[8] Alexander Cowtan, Silas Dilkes, Ross Duncan, Alexandre Krajenbrink, Will Simmons, and Seyon Sivarajah, 'On the qubit routing problem', *arXiv preprint arXiv:1902.08091*, (2019).

[9] Burkard Rainer Ernst, Eranda Dragoti-Cela, P.M. Pardalos, and L.S. Pitsoulis, *The quadratic assignment problem*, volume 2, 241–337, 1998.

[10] P. Eyerich, R. Mattmüller, and G. Röger, 'Using the context-enhanced additive heuristic for temporal and numeric planning', in *Proceedings of the* 19^{th} *International Conference on Automated Planning and Scheduling*, pp. 318 – 325, (2009).

[11] E. Farhi, J. Goldstone, and S. Gutmann., 'A quantum approximate optimization algorithm.', in *arXiv preprint arXiv:1411.4028*, (2014).

[12] A. Gerevini, L. Saetti, and I. Serina, 'Planning through stochastic local search and temporal action graphs', *Journal of Artificial Intelligence Research*, **20**, 239 – 290, (2003).

[13] Pranav Gokhale, Yongshan Ding, Thomas Propson, Christopher Winkler, Nelson Leung, Yunong Shi, David I Schuster, Henry Hoffmann, and Frederic T Chong, 'Partial compilation of variational algorithms for noisy intermediate-scale quantum machines', in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, pp. 266–278. ACM, (2019).

[14] Stuart Hadfield, Zhihui Wang, Bryan O'Gorman, Eleanor G Rieffel, Davide Venturelli, and Rupak Biswas, 'From the quantum approximate optimization algorithm to a quantum alternating operator ansatz', *Algorithms*, **12**(2), 34, (2019).

[15] Malte Helmert, 'The fast downward planning system', *Journal of Artificial Intelligence Research*, **26**, 191246, (2006).

[16] Toshinari Itoko, Rudy Raymond, Takashi Imamichi, and Atsushi Matsuo, 'Optimization of quantum circuit mapping using gate transformation and commutation', *Integration*, (2019).

[17] Gushu Li, Yufei Ding, and Yuan Xie, 'Tackling the qubit mapping problem for nisq-era quantum devices', in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pp. 1001–1014. ACM, (2019).

[18] Prakash Murali, Norbert Matthias Linke, Margaret Martonosi, Ali Javadi Abhari, Nhung Hong Nguyen, and Cinthia Huerta Alderete, 'Full-stack, real-system quantum computer studies: Architectural comparisons and design insights', *arXiv preprint arXiv:1905.11349*, (2019).

[19] Angelo Oddi and Riccardo Rasconi, 'Greedy randomized search for scalable compilation of quantum circuits', in *Proceedings of the Fifteenth International Conference on the Integration of Constraint Programming, Artificial Intelligence, and Operations Research (CPAIOR2018)*, pp. 446–461, (2018).

[20] Bryan O'Gorman, William J. Huggins, Eleanor G. Rieffel, and K. Birgitta Whaley. Generalized swap networks for near-term quantum computing, 2019.

[21] Riccardo Rasconi and Angelo Oddi, 'An innovative genetic algorithm for the quantum circuit compilation problem', in *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI2019)*, p. In press, (2019).

[22] Davide Venturelli, Minh Do, Bryan O'Gorman, Jeremy Frank, Eleanor Rieffel, Kyle EC Booth, Thanh Nguyen, Parvathi Narayan, and Sasha Nanda, 'Quantum circuit compilation: An emerging application for automated reasoning', *ICAPS 2019 Workshop SPARK*, (2019).

[23] Davide Venturelli, Minh Do, Eleanor Rieffel, and Jeremy Frank, 'Temporal planning for compilation of quantum approximate optimization circuits', in *Proceedings of The 26th International Joint Conference on Artificial Intelligence (IJCAI17)*, (2017).

[24] Vincent Vidal and Hctor Geffner, 'Branching and pruning: An optimal temporal pocl planner based on constraint programming', *Artificial Intelligence Journal*, **170(3)**, 298335, (2006).

[25] Zhihui Wang, Nicholas C Rubin, Jason M Dominy, and Eleanor G Rieffel, 'xy-mixers: analytical and numerical results for qaoa', *arXiv preprint arXiv:1904.09314*, (2019).

[26] Alwin Zulehner, Alexandru Paler, and Robert Wille, 'An efficient methodology for mapping quantum circuits to the ibm qx architectures', *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, (2018).