

Mapping Quantum Circuits to IBM QX Architectures Using the Minimal Number of SWAP and H Operations

Robert Wille

Institute for Integrated Circuits, Johannes Kepler University Linz, Austria
robert.wille@jku.at

Lukas Burgholzer

Institute for Integrated Circuits, Johannes Kepler University Linz, Austria
lukas.burgholzer@jku.at

Alwin Zulehner

Institute for Integrated Circuits, Johannes Kepler University Linz, Austria
alwin.zulehner@jku.at

ABSTRACT

The recent progress in the physical realization of quantum computers (the first publicly available ones—IBM’s QX architectures—have been launched in 2017) has motivated research on automatic methods that aid users in running quantum circuits on them. Here, certain physical constraints given by the architectures which restrict the allowed interactions of the involved qubits have to be satisfied. Thus far, this has been addressed by inserting SWAP and H operations. However, it remains unknown whether existing methods add a minimum number of SWAP and H operations or, if not, how far they are away from that minimum—an \mathcal{NP} -complete problem. In this work, we address this by formulating the mapping task as a symbolic optimization problem that is solved using reasoning engines like Boolean satisfiability solvers. By this, we do not only provide a method that maps quantum circuits to IBM’s QX architectures with a *minimal* number of SWAP and H operations, but also show by experimental evaluation that the number of operations added by IBM’s heuristic solution exceeds the lower bound by more than 100% on average. An implementation of the proposed methodology is publicly available at http://iic.jku.at/eda/research/ibm_qx_mapping.

1 INTRODUCTION

Quantum computing [15] currently gains considerable momentum. First concepts showing the (theoretical superiority) of this new computation paradigm have already been developed a few decades ago (e.g., by Shor’s famous factorization algorithm proposed in [17] or Grover’s database search iteration proposed in [7]). While this mainly electrified the academic community only, a real “push” of this topic was caused by the involvement of “big players” such as IBM, Google, Intel, or Microsoft in the recent years—additionally triggering the interest of industry as well as the public at large [3, 6, 8]. They are aiming to utilize quantum computers for applications such as quantum chemistry, optimization, machine learning, cryptography, quantum simulation, or solving systems of linear equations [16].

Here, particularly IBM’s approach stands out, which provided the first publicly available quantum processor (in 2017) that can be accessed by everyone (not only academics) through cloud access. Since then, their machines have been used by more than 100,000 users, who have run more than 6.5 million experiments thus far [9]. This motivated research on automatic methods that aid users in running quantum circuits on the corresponding machines (known as *IBM QX architectures*).

An obvious problem is thereby how to efficiently realize the desired quantum functionality on a respectively given architecture. The functionality is usually represented by means of a *quantum circuit* composed of qubits and quantum operations. *Qubits* provide the basic entity of quantum computers which, as in conventional computation, may assume two basis states $|0\rangle$ and $|1\rangle$ but additionally also any superpositions of them. *Quantum operations* can be defined by arbitrary unitary matrices but have to be decomposed into elementary operations which are supported by the given architecture. While for decomposing arbitrary quantum functionality to a sequence of elementary operations many solutions already exist (by specifying the decomposition manually [4] or using automated approaches as described in [1, 14]), further constraints need to be addressed.

In fact, logical qubits used in the originally given quantum circuit description cannot arbitrarily be mapped to physical qubits used in the QX architectures, but have to satisfy certain constraints defined by a coupling map. This can be accomplished by adding *SWAP* and *H operations*—increasing the size of the circuit and, by this, harming the fidelity of the execution. Accordingly, this raises the question of how to derive a proper mapping of logical qubits to physical qubits while, at the same time, minimizing the number of added SWAP and H operations—an \mathcal{NP} -complete problem as recently proven in [2]. In the recent past, several solutions for this have been proposed [12, 13, 18, 22, 23]. Moreover, even competitions seeking the best possible solution for this problem have been conducted in order to further trigger development in this area (see [11]).

However, none of the methods presented thus far solves this problem in an *exact*, i.e., a minimal, fashion. Instead, heuristics are applied. While this is a reasonable strategy for an \mathcal{NP} -complete problem, it leads to a significant uncertainty about the quality of the developments outlined above: Even if significant progress can be observed, it still remains completely unclear how far the proposed heuristics are from the optimum and, hence, how good the proposed methods really are? Because of this, *exact* methods that can generate minimal (or at least close-to-minimal) results are essential for a substantial evaluation of the current state of the art—even if the optimum can only be generated for small instances.

In this work, we are proposing such solutions. To this end, we consider *all* possible applications of the SWAP and H operations that may influence the realization of an originally given circuit on a QX architecture—a computationally very expensive task. In order to cope with this complexity, we propose to utilize powerful reasoning engines such as solvers for Boolean satisfiability that can cope with large search spaces. Besides that, additional performance optimizations are proposed that may lead to solutions which are not guaranteed to be minimal anymore, but can significantly speed up the solving time while still generating at least close-to-minimal solutions. This is confirmed by experimental evaluations, which additionally show that IBM’s heuristic mapping solution exceeds the lower bound by more than 100% on average.

2 BACKGROUND

To keep this paper self contained, we briefly review quantum circuits as well as IBM’s QX architectures in this section.

2.1 Quantum Circuits

Quantum circuits are a frequently used description means for quantum computations. Here, the logical qubits are represented by vertical circuit lines, while quantum gates describe (from left to right) the order in which quantum operations (whose functionality is described by unitary matrices) are applied to the qubits. Common operations that act on a single qubit are X , H , and T , which negate the state of a qubit, set it into superposition, or apply a phase shift, respectively.¹ In a quantum circuit diagram, these operations are denoted by square boxes labeled with a corresponding identifier.

Besides operations working on a single qubit, there also exist operations acting on multiple ones. A subset of such operations are controlled operations. Here, a single qubit operation is applied to the target qubit if all controlling qubits are in the basis state $|1\rangle$.

¹Since the actual functionality of the single quantum operations are not relevant in this work, we refer, e.g., to [15] for a more detailed treatment on that.




Figure 1: Quantum circuit to be mapped

One commonly used representative is the CNOT gate, where a single control qubit determines whether the state of the target qubit is inverted or not. The CNOT gate plays an important role in quantum computations, since it—together with single qubit operations—provides a universal set for quantum computing, i.e., any quantum computation can be decomposed into a sequence of CNOT and single qubit gates.

In the remainder of this work, we are using the following notation for quantum circuits:

DEFINITION 1. Let $Q = \{q_1, q_2, \dots, q_j, \dots, q_n\}$ be a set of n logical qubits. Then, a quantum circuit $G = g_1 g_2 \dots g_k \dots g_{|G|}$ is a sequence of quantum gates, where each gate g_k is either a single qubit gate $U_k(q_j, U)$ with target qubit $q_j \in Q$ and a unitary matrix U describing the corresponding functionality, or a controlled not gate $CNOT_k(q_c, q_t)$ with control qubit $q_c \in Q$ and target qubit $q_t \in Q$ (with $q_c \neq q_t$).

EXAMPLE 1. Fig. 1a shows a quantum circuit composed of 4 qubits and 8 gates. The single qubit gates H and T are visualized by square boxes labeled with H and T , respectively, while the control and target qubit of CNOT gates are denoted by \bullet and \oplus , respectively.

2.2 IBM QX Architectures

IBM's QX architectures have been made publicly available through cloud access in 2017 to allow conducting quantum computations on real devices. The IBM QX architectures provide the universal single qubit gate $U(\theta, \phi, \lambda) = R_z(\phi)R_y(\theta)R_z(\lambda)$ that is composed of two rotations around the z -axis and a rotation around the y -axis. Since the U gate is universal, any single qubit operation can be conducted by specifying the parameters θ, ϕ, λ . By additionally supporting CNOT gates, IBM QX architectures allow for universal quantum computing (even though still limited in the number of qubits and gate fidelity).

To run quantum circuits on IBM's QX architectures, the respective logical qubits have to be mapped to physical ones and gates have to be decomposed into a sequence of U and $CNOT$ gates. In this work, we assume that the decomposition step has already been conducted (by specifying the decomposition manually [4] or using automated approaches as described in [1, 14]). However, even then further constraints given by the architectures (called *CNOT-constraints*) have to be satisfied. They state that not all physical qubits can interact with each other and, hence, CNOT gates cannot be applied to arbitrary pairs of physical qubits. Moreover, even if a CNOT gate can be applied to two physical qubits, it is restricted which physical qubit may serve as control and which physical qubit as target. These constraints are defined in a coupling map.

DEFINITION 2. Let $P = \{p_1, p_2, \dots, p_i, \dots, p_m\}$ be a set of m physical qubits available in the architecture. Then, the coupling map $CM \subseteq P \times P$ defines which physical qubits can interact with each other. More precisely, $(p_i, p_j) \in CM$ states that a CNOT gate with control qubit p_i and target qubit p_j can be applied, while $(p_i, p_j) \notin CM$ states that a CNOT gate with control qubit p_i and target qubit p_j can not be applied.

EXAMPLE 2. Figure 2 graphically represents the coupling map $CM = \{(p_2, p_1), (p_3, p_1), (p_3, p_2), (p_4, p_3), (p_4, p_5), (p_5, p_3)\}$ for the




Figure 2: Coupling map of IBM QX4 [10]




Figure 3: Decomposition of a SWAP operation

IBM QX4 architecture as directed graph. More precisely, the physical qubits p_1, p_2, \dots, p_5 are visualized as vertices, whereas an entry $(p_i, p_j) \in CM$ is visualized as directed edge from p_i to p_j . Hence, an arrow pointing from p_i to p_j indicates that a CNOT gate with control qubit p_i and target qubit p_j can be applied.

To satisfy the CNOT-constraints, one has to map the n logical qubits q_1, q_2, \dots, q_n of the circuit to the $m \geq n$ physical qubits p_1, p_2, \dots, p_m of the considered quantum device such that all constraints given by the corresponding coupling map are satisfied. Unfortunately, it is usually not possible to find a mapping such that the constraints are satisfied throughout the whole circuit. More precisely, the following problems may occur:

- A CNOT gate $CNOT_k(q_c, q_t)$ shall be applied while q_c and q_t are mapped to physical qubits p_i and p_j , respectively, and $(p_i, p_j) \notin CM$ as well as $(p_j, p_i) \notin CM$.
- A CNOT gate $CNOT_k(q_c, q_t)$ shall be applied while q_c and q_t are mapped to physical qubits p_i and p_j , respectively, and $(p_i, p_j) \notin CM$ while $(p_j, p_i) \in CM$.

To overcome these problems, one strategy is to insert additional gates into the circuit to be mapped. More precisely, to overcome the first issue, one can insert SWAP operations into the circuit that exchange the state of two physical qubits and, by this, move around the logical ones (a strategy which also has been applied to make quantum circuits nearest neighbor-compliant [21]).

EXAMPLE 3. Fig. 3 shows the effect of a SWAP gate as well as its decomposition into elementary gates supported by the QX architectures. Assume that the logical qubits q_1 and q_2 are initially mapped to the physical ones p_1 and p_2 , respectively (indicated by \leftarrow). Then, by applying a SWAP gate, the states of p_1 and p_2 are exchanged—eventually yielding a mapping where q_1 and q_2 are mapped to p_2 and p_1 , respectively.

The second issue may also be solved by inserting SWAP operations. However, it is cheaper (fewer overhead is generated) to insert four Hadamard operations (labeled by H) as they switch the direction of the CNOT gate (i.e., to change the target and the control qubit). This can also be observed in Fig. 3, where H gates switch the direction of the middle CNOT in order to satisfy all CNOT-constraints given by the coupling map.

As cost metric, we utilize the number of operations, since each operation introduces an error with a certain probability. Therefore, inserting a SWAP operation increases the cost by 7 (cf. Fig. 3), whereas switching the direction of a CNOT gate increases the cost by 4 (since 4 H gates are added). Obviously, the general objective is to keep the overall cost as low as possible to keep the overall fidelity as high as possible.

3 DETERMINING A MINIMAL SOLUTION

Recently, first solutions have been proposed which derive a proper mapping of logical qubits to physical qubits while, at the same time, satisfying the CNOT constraints (see, e.g., [12, 13, 18, 22, 23]).

However, as discussed in Section 1, it remains unknown how well they address the general objective reviewed in the previous section (i.e., how well they keep the costs caused by adding SWAP and H operations as small as possible). In this section, we describe how to determine a minimal solution for this problem. To this end, we first describe the main idea of tackling the underlying complexity of the problem and, afterwards, introduce a symbolic formulation describing the problem in terms of a Boolean function. This is eventually used to solve the problem by applying powerful and efficient reasoning engines.

3.1 Main Idea for Tackling the Complexity

In this work, we aim for determining minimal or at least close-to-minimal solutions for mapping quantum circuits to IBM QX architectures.² To this end, we cannot heuristically or incompletely consider the search space but have to consider *all* possible applications of the SWAP and H operations that may influence the realization of an originally given circuit on a QX architecture. Obviously, this results in a computationally very expensive task (\mathcal{NP} -complete as recently proven in [2]). In order to cope with this complexity, we propose to utilize powerful reasoning engines such as solvers for Boolean satisfiability that can deal with large search spaces. These solvers address the satisfiability problem defined as follows:

DEFINITION 3. *The satisfiability problem determines an assignment to the variables of a Boolean function $\Phi : \{0, 1\}^n \rightarrow \{0, 1\}$ such that Φ evaluates to 1 or proves that no such assignment exists. In an extended interpretation, additionally an objective function \mathcal{F} defined by $\mathcal{F}(x_1, \dots, x_n) = \sum_{i=1}^n w_i x_i$ with $w_1, \dots, w_n \in \mathbb{Z}$ and $\dot{x} \in \{\bar{x}_i, x_i\}$ is provided. In this case, an assignment is to be determined which does not only satisfy Φ but, at the same time, minimizes \mathcal{F} .*

EXAMPLE 4. *Let $\Phi = (x_1 + x_2 + \bar{x}_3)(\bar{x}_1 + x_3)(\bar{x}_2 + x_3)$. Then, $x_1 = 1$, $x_2 = 0$, and $x_3 = 1$ is a satisfying assignment solving the SAT problem. Additionally, let $\mathcal{F} = x_1 + x_2 + x_3$. Then, $x_1 = 0$, $x_2 = 0$, and $x_3 = 0$ is a solution which does not only satisfies Φ but also minimizes \mathcal{F} .*

In the past, very efficient reasoning engines for the satisfiability problem have been proposed (see, e.g., [5, 19]). Instead of simply traversing the complete space of assignments, intelligent decision heuristics, powerful learning schemes, and efficient implication methods are applied and, by this, instances composed of numerous variables and constraints—defining huge search spaces—can be solved efficiently. In this work, we are utilizing this reasoning power to consider the whole search space and, by this, to tackle the complexity. To this end, however, a symbolic formulation is required which completely describes the problem and is provided in terms of a Boolean function (so that it can be used by those reasoning engines).

3.2 Symbolic Formulation of the Problem

In the following, a symbolic formulation of the considered problem—mapping quantum circuits to IBM QX architectures using SWAP and H operations—is proposed. To this end, variables are defined which describe all possible applications of the SWAP and H operations. More precisely, those operations basically affect how logical qubits from an originally given quantum circuit are mapped to the physical qubits of an IBM QX architecture. The mapping might be changed before each gate. Since only CNOT gates may cause violations of CNOT constraints, we ignore single qubit gates when

formulating the mapping problem.³ This leads to the following symbolic formulation:

DEFINITION 4. *Let $G = g_1 g_2 \dots g_k \dots g_{|G|}$ be a quantum circuit composed of $|G|$ CNOT gates. Each gate g_k operates on a (logical) control qubit q_c and a (logical) target qubit q_t (cf. Def. 1). Furthermore, let $Q = \{q_1, \dots, q_j, \dots, q_n\}$ be a set of n logical qubits that shall be mapped to the $m \geq n$ physical qubits in $P = \{p_1, \dots, p_i, \dots, p_m\}$. Finally, let $CM \subseteq P \times P$ be the description of the coupling map indicating what circuit lines can interact with each other on the given architecture and how (cf. Def. 2). Then, mapping variables x_{ij}^k with $k \in \{1, \dots, |G|\}$, $i \in \{1, \dots, m\}$, and $j \in \{1, \dots, n\}$ are introduced representing whether, before gate $g_k \in G$, the logical qubit q_j is mapped to the physical qubit p_i ($x_{ij}^k = 1$) or not ($x_{ij}^k = 0$).*

EXAMPLE 5. *Consider again the circuit shown in Fig. 1a (and assume the single qubit gates have been removed as shown in Fig. 1b). Then, Fig. 4 sketches a symbolic formulation for mapping the circuit to IBM QX4 which represents all possible mappings of logical qubits to physical qubits. For example, the leftmost part of Fig. 4 represents the initial mapping of the logical qubits to the physical ones. Here, e.g., setting $x_{13}^1 = 1$ represents that logical qubit q_3 is mapped to physical qubit p_1 right before gate g_1 .*

Passing this symbolic formulation to a reasoning engine would yield arbitrary assignments that most likely encode impossible/useless mappings (e.g., mapping several logical qubits to the same physical one). Hence, we have to restrict the assignment of variables so that only valid solutions are obtained. To this end, we have to ensure that:

- (1) A well-defined mapping between logical and physical qubits is conducted (i.e., each logical qubit is uniquely assigned to exactly one physical qubit and vice versa). This is ensured by

$$\bigwedge_{k=1}^{|G|} \left(\bigwedge_{j=1}^n \left(\sum_{i=1}^m x_{ij}^k = 1 \right) \wedge \bigwedge_{i=1}^m \left(\sum_{j=1}^n x_{ij}^k \leq 1 \right) \right). \quad (1)$$

- (2) All gates only act on physical qubits that have a corresponding entry in the coupling map of the considered architecture or an entry where control and target qubit are switched.⁴ This is ensured by

$$\bigwedge_{g_k = \text{CNOT}_{(q_c, q_t)} \in G} \left(\bigvee_{(p_i, p_j) \in CM} (x_{ic}^k \wedge x_{jt}^k) \vee (x_{it}^k \wedge x_{jc}^k) \right). \quad (2)$$

Adding these restrictions and passing the resulting symbolic formulation to a reasoning engine eventually yields a valid solution. Moreover, the resulting formulation covers the entire search space in a symbolic fashion. Having this, all that is left is a proper description of the costs of the respectively chosen mapping. As reviewed in Section 2.2, these costs accumulate from (1) the costs for changing the mapping of the logical qubits to the physical ones throughout the circuit (by inserting SWAP operations) and (2) the costs for switching the control and the target qubit for CNOT gates (by inserting H operations). The former one (changes on the mapping) may be applied before each CNOT gate (except the first one, which defines the initial mapping and, hence, can be set arbitrarily anyway); the latter one (switching the control/target qubits) may happen in each gate. To properly describe this within the symbolic formulation, we further introduce the following variables:

²As discussed in Section 2.2, we solve the mapping problem by inserting SWAP and/or H operations and assume that quantum circuits are already decomposed into elementary operations. Moreover, we do not consider pre- or post-mapping optimizations (as, e.g., proposed in [12, 23]) that may be applied before/after the mapping, but solely consider the actual mapping process in an exact fashion.

³Note that this additionally reduces the overall complexity of the problem to be solved.

⁴Note that we additionally consider entries where the control and target qubit are switched since this can be handled by inserting H gates. This way, the reasoning engines get to decide what combination of SWAP and H gates yield the cheapest global mapping.




Figure 4: Symbolic formulation for mapping the circuit shown in Fig. 1a

DEFINITION 5. Let $1 \leq k \leq |G|$ be the index of gate g_k in a quantum circuit, m the number of physical qubits in the considered quantum device, and $\pi \in \Pi$ a permutation of m elements that indicates how the state of the physical qubits is permuted (eventually realized by inserting SWAP operations). Then, the permutation variables y_π^k indicate whether the permutation π is applied before gate g_k ($y_\pi^k = 1$) or not ($y_\pi^k = 0$). Furthermore, the switching variables z^k indicate whether the direction of the CNOT gate g_k is switched ($z^k = 1$) or not ($z^k = 0$).

EXAMPLE 6. Consider again the symbolic formulation shown in Fig. 4. The spots in the circuit where the mapping may change are sketched by boxes labeled π . Here, the variable assignment before and after (i.e., the assignments of x_{ij}^{k-1} and x_{ij}^k) may change according to a permutation $\pi \in \Pi$ (eventually to be represented by y_π^k). Furthermore, in each gate g_k , the z^k -variables define whether the direction of the CNOT gate is switched or not.

Using these variables, we can describe what permutation $\pi \in \Pi$ is applied to the states of the physical qubits of the circuit lines before each gate $g_k \in G$ by introducing

$$\bigwedge_{k=2}^{|G|} \left(\bigwedge_{\pi \in \Pi} \left(\bigwedge_{i=1}^m \bigwedge_{j=1}^m (x_{ij}^{k-1} = x_{\pi(i)j}^k) \right) \Leftrightarrow y_\pi^k \right). \quad (3)$$

In fact, this ensures that y_π^k is set to 1 iff the assignment of the variables x_{ij}^{k-1} and x_{ij}^k indeed describe a change of the mapping defined by π .⁵

Similarly, we can describe for each CNOT gate g_k whether the direction of the control and target qubits are switched by introducing

$$g_k = \text{CNOT}_{T_c, q_t} \in G \left(\bigvee_{(p_i, p_j) \in CM} (x_{it}^k \wedge x_{jc}^k) \right) \Leftrightarrow z^k. \quad (4)$$

In fact, this ensures that z^k is set to 1 iff, for gate g_k , the control qubit is set to position j and the target qubit is set to position i although, according to the coupling map, it has to be vice versa (i.e., control and target qubits are switched).

Satisfying all of the constraints introduced above yields a valid mapping of the originally given circuit to the desired architecture while, at the same time, the costs are determined by

$$\mathcal{F} = \sum_{k=2}^{|G|} \sum_{\pi \in \Pi} (7 \cdot \text{swaps}(\pi) y_\pi^k) + \sum_{k=1}^{|G|} (4 \cdot z^k). \quad (5)$$

Here, $\text{swaps}(\pi)$ defines the number of SWAP operations needed to realize the permutation π . This has to be determined for each permutation $\pi \in \Pi$ —a process, which needs to be conducted only once and can be done, e.g., by using an exhaustive search for the architectures

⁵If $n < m - 1$, π cannot be determined uniquely. Then, a left-handed implication is required instead of an equivalence in Eq. (3)—in conjunction with a constraint that only one variable y_π^k is assigned to 1. For sake of clarity of Eq. (3), we assume that $n = m$.




Figure 5: Resulting circuit (with minimal SWAP/H costs)

considered in this work. By this, whenever the reasoning engine chooses a mapping which eventually creates a permutation π before gate g_k , Eq. (3) sets $y_\pi^k = 1$ and, hence, adds the corresponding costs (7 gates for each SWAP operation; cf. Section 2) to the overall costs \mathcal{F} . Similarly, whenever the reasoning engine chooses a mapping which requires switching the direction of a CNOT gate g_k , Eq. (4) sets $z^k = 1$ and, hence, adds the corresponding costs (4 H operations; cf. Section 2) to the overall costs \mathcal{F} .

3.3 Minimizing the Cost

Passing the eventually resulting symbolic formulation to a reasoning engine allows to determine a valid mapping together with the associated cost (i.e., the number of additionally required elementary operations). Since we are also interested in the minimum costs, the cost function \mathcal{F} needs to be further restricted. One direct solution could be to simply set \mathcal{F} to a fixed value and approach towards the minimum, e.g., by applying a binary search. However, since many reasoning engines additionally allow to consider an objective function (cf. Def. 3), the most efficient way is to simply add the objective $\min : \mathcal{F}$ to the resulting instance—enforcing the reasoning engine not only to determine a satisfying assignment (representing a valid mapping) but, at the same time, also to minimize \mathcal{F} .

EXAMPLE 7. Passing the symbolic formulation sketched in Fig. 4 together with all constraints and the objective function to a reasoning engine, eventually yields a mapping (and a corresponding addition of SWAP and H operations) as shown in Fig. 5. This circuit provides a realization of the originally given circuit from Fig. 1a which is applicable for the IBM QX architecture specified by the coupling map shown in Fig. 2 and, at the same time, yields minimum costs caused by additionally required SWAP and H operations ($\mathcal{F} = 4$).

4 PERFORMANCE IMPROVEMENTS

Determining minimal solutions obviously is the desired way to go. However, even with powerful reasoning engines, we cannot always escape the \mathcal{NP} -complete complexity of the problem. In this regard, the methodology proposed in the previous section allows for several performance improvements. In fact, the reasoning engine only needs to determine a “minimal” assignment for the x_{ij}^k -variables (the variables y_π^k and z^k can be ignored, since they are only used for formulating the costs and their assignments can directly be deduced from the x_{ij}^k -variables). For a quantum circuit composed of n logical qubits and $|G|$ CNOT gates to an architecture with m physical qubits, this leads to a total of $n \cdot m \cdot |G|$ variables to be assigned and, hence, an overall search space of $2^{n \cdot m \cdot |G|}$ which can easily be restricted by adding further constraints to the x_{ij}^k -variables. While

this may lead to solutions which are not guaranteed to be minimal anymore, it can significantly speed up the solving time while, at the same time, remaining very close-to-minimal (as also confirmed by experimental evaluations summarized in Section 5). This section shows possible improvements in this regard.

4.1 Considering Subsets of Physical Qubits

A scenario frequently occurs where the number n of logical qubits of a given quantum circuit to be mapped is smaller than the number m of physical qubits provided by the architecture (i.e., where $n < m$). Then, obviously, not all physical qubits are required. In fact, this allows to consider only a subset of n physical qubits while ignoring the remaining $m - n$ ones. Since the number of physical qubits to consider contributes to the search space in an exponential fashion, restricting this number yields substantial simplifications. In order to remain as close as possible to the minimal solution, one can try out all $\binom{m}{n}$ possible subsets of qubits to consider and solve the respectively resulting (smaller) instances separately. This reduces the overall search space to $\binom{m}{n}2^{n^2 \cdot |G|}$.

EXAMPLE 8. Consider again the symbolic instance for mapping a four qubit quantum circuit to a five qubit architecture sketched in Fig. 4. By considering only four physical qubits in the mapping procedure, the overall search space for a single instance reduces from $2^{4 \cdot 5 \cdot 5} = 2^{100}$ to $2^{4^2 \cdot 5} = 2^{80}$. Even if all $\binom{5}{4} = 5$ possible subsets of physical qubits are considered separately, this still yields a significant reduction of the overall search space.

The search space can be reduced further by checking whether some of the physical qubits in a subset are isolated from others (this can be done in $O(n)$ time). If so, the instance for this subset does not have to be passed to the reasoning engine as no solution can be found anyway.

EXAMPLE 9. Assume that the circuit over four qubits considered thus far shall be mapped to the IBM QX4 architecture shown in Fig. 2. Then, all subsets of physical qubits that have to be checked should contain p_3 , since no connected sub-graph composed of four nodes without p_3 is possible. This reduces the number of instances that are passed to the reasoning engine from $\binom{5}{4} = 5$ to $\binom{4}{3} = 4$.

4.2 Restricting the Possible Permutations

Thus far, we allowed permutations $\pi \in \Pi$ of the mapping before each gate (except the first one where an arbitrary initialization can be chosen). While this guarantees minimality (since all possible solutions are considered), this substantially contributes to the complexity. In many cases, however, valid and cheap mappings are still possible if permutations of mappings are allowed not before all gates $g \in G$, but only before a subset $G' \subseteq G \setminus \{g_1\}$ of them. With $|G'|$ being significantly smaller than G , this reduces the overall search space to $2^{n \cdot m \cdot (|G'|+1)}$. While applying this idea, G' can be chosen arbitrarily. A smaller G' leads to a larger performance improvement, but also a more restricted instance (yielding solutions that might be far from minimal or even instances for which no valid mapping can be determined anymore). In this work, the following strategies for defining G' are considered:⁶

- *Disjoint qubits*, which exploits the fact that gates acting on disjoint sets of qubits can always be mapped in a way that no intermediate permutations are required.⁷ To this end, the quantum circuit is clustered into sequences of gates acting on disjoint sets of qubits and permutations are only allowed before each of those sequences.

- *Odd gates*, which allows permutations only before gates with an odd index (except for g_1). Here, it is still guaranteed that a valid mapping can be determined since either (1) the gates operate on the disjoint sets of qubits as discussed above, (2) the gates share both qubits, or (3) the gates share one qubit (and there exists at least one qubit that can interact with two other qubits).
- *Qubit triangle*, which exploits the structure of architectures whose coupling map forms “triangles” of physical qubits as in case of, e.g., p_1, p_2 , and p_3 in Fig. 2. Here, we can cluster the circuit into sequences of gates where each sequence acts on at most three qubits. Then, each such sequence of gates can be mapped to a triangle as described above and permutations are only required before each of those sequences.

EXAMPLE 10. Consider the quantum circuit shown in Fig. 1b. Applying the strategies proposed above yields the following subsets G' :

- Disjoint qubits: $G' = \{g_3, g_4, g_5\}$, since the gates g_1 and g_2 operate on disjoint qubits (saving the permutation between g_1 and g_2).
- Odd gates: $G' = \{g_3, g_5\}$, since they constitute the odd gates in the circuit.
- Qubit triangle: $G' = \{g_2\}$, since all gates g_2, g_3, g_4 , and g_5 operate on only three qubits and, hence, can be mapped to one of the “triangles” of the architecture without the need for further permutations. That is, only a permutation prior to g_2 needs to be considered.

As can be seen, these strategies yield much more restrictive applications of permutations. While this substantially increases the performance of the solving process, it does not harm minimality in this case (but may for other circuits as evaluated in the next section).

5 EXPERIMENTAL RESULTS

The proposed method for mapping a quantum circuit to the IBM QX architectures using the minimal number of SWAP and H operations has been implemented in C++ (the implementation is publicly available at http://iic.jku.at/eda/research/ibm_qx_mapping). As reasoning engine, the Z3 solver [5] has been utilized. Afterwards, we conducted extensive evaluations using quantum circuits (also considered in previous work and taken from [4, 20]) to be mapped to the IBM QX4 architecture [10]. All evaluations have been conducted on an Intel Core i7-3930K machine with 4 GHz and 64 GB of main memory running Ubuntu 18.04.

The left part of Table 1 thereby provides a selection of the results obtained by a first series of evaluations—aiming for evaluating the effect of the performance improvements discussed in Section 4. More precisely, the first three columns describe the name of the considered quantum circuit, the number of logical qubits n , and the *original cost* of the circuit (i.e., the number of single qubit gates plus the number of CNOT gates before mapping). In the remaining columns, we list the cost c (i.e., the number of gates) of the obtained circuit and the runtime t (in CPU seconds) required by Z3 when applying the method guaranteeing minimality discussed in Section 3 as well as the adapted methods additionally incorporating the performance improvements discussed in Section 4. For the adapted versions, we additionally list the difference to the minimum (i.e., Δ_{min}) in parenthesis. Moreover, for each strategy that limits the number of permutations, we additionally list in columns denoted $|G'|$ how many permutations are allowed.

As can be seen by these results, determining a mapping with the minimum number of SWAP and H gates is quite expensive—only solutions for instances with rather few CNOT gates can be determined (which is not surprising since the underlying problem is \mathcal{NP} -complete). When considering only a subset of the physical qubits (cf. Section 4.1), we observe a significant reduction of the runtime for benchmarks with 3 or 4 qubits, while still preserving

⁶Many more strategies have been considered and evaluated, but are omitted in this paper due to space limitations.

⁷Note that such a set of gates is called *layer* in some heuristic solutions [12, 22].

Table 1: Experimental results

Benchmark	n	original cost	Min. (Sec. 3)		Perf. Opt. (Sec. 4.1)		Performance Optimized (Section 4.2)								IBM [12]	
			c_{min}	t [s]	c (Δ_{min})	t [s]	Disjoint qubits		Odd gates		Qubit triangle					
							$ G' $	c (Δ_{min})	t [s]	$ G' $	c (Δ_{min})	t [s]	$ G' $	c (Δ_{min})	t [s]	c (Δ_{min})
3_17_13	3	19 + 17 = 36	59	29	59 (+0)	0	17	59 (+0)	0	9	60 (+1)	0	1	60 (+1)	0	80 (+21)
ex-1_166	3	10 + 9 = 19	31	5	31 (+0)	0	9	31 (+0)	0	5	31 (+0)	0	1	31 (+0)	0	39 (+8)
ham3_102	3	9 + 11 = 20	36	10	36 (+0)	0	11	36 (+0)	0	6	36 (+0)	0	1	36 (+0)	0	48 (+12)
miller_11	3	27 + 23 = 50	82	231	82 (+0)	0	23	82 (+0)	0	12	82 (+0)	0	1	82 (+0)	0	82 (+0)
4gt11_84	4	9 + 9 = 18	34	7	34 (+0)	0	9	34 (+0)	0	5	34 (+0)	0	2	34 (+0)	0	37 (+3)
rd32-v0_66	4	18 + 16 = 34	63	281	63 (+0)	35	16	63 (+0)	35	8	63 (+0)	1	2	72 (+9)	0	101 (+38)
rd32-v1_68	4	20 + 16 = 36	65	276	65 (+0)	35	16	65 (+0)	36	8	65 (+0)	1	2	74 (+9)	0	99 (+34)
4gt11_82	5	9 + 18 = 27	62	133	62 (+0)	137	18	62 (+0)	139	9	62 (+0)	3	5	62 (+0)	1	77 (+15)
4gt11_83	5	9 + 14 = 23	49	17	49 (+0)	17	14	49 (+0)	18	7	50 (+1)	1	3	50 (+1)	0	65 (+16)
4gt13_92	5	36 + 30 = 66	109	528	109 (+0)	533	29	109 (+0)	199	15	110 (+1)	10	9	110 (+1)	5	126 (+17)
4mod5-v0_19	5	19 + 16 = 35	64	256	64 (+0)	264	16	64 (+0)	255	8	68 (+4)	2	3	69 (+5)	0	109 (+45)
4mod5-v0_20	5	10 + 10 = 20	35	10	35 (+0)	10	10	35 (+0)	11	5	35 (+0)	0	3	35 (+0)	0	64 (+29)
4mod5-v1_22	5	10 + 11 = 21	40	7	40 (+0)	7	10	40 (+0)	9	6	40 (+0)	0	3	43 (+3)	0	52 (+12)
4mod5-v1_24	5	20 + 16 = 36	63	54	63 (+0)	55	16	63 (+0)	56	8	63 (+0)	3	3	63 (+0)	0	98 (+35)
alu-v0_27	5	19 + 17 = 36	63	74	63 (+0)	73	16	63 (+0)	38	9	63 (+0)	2	3	67 (+4)	0	101 (+38)
alu-v1_28	5	19 + 18 = 37	64	94	64 (+0)	92	17	64 (+0)	44	9	67 (+3)	10	3	68 (+4)	0	123 (+59)
alu-v1_29	5	20 + 17 = 37	64	351	64 (+0)	355	16	64 (+0)	119	9	64 (+0)	3	3	68 (+4)	0	104 (+40)
alu-v2_33	5	20 + 17 = 37	64	42	64 (+0)	44	17	64 (+0)	44	9	64 (+0)	4	4	64 (+0)	0	99 (+35)
alu-v3_34	5	28 + 24 = 52	90	719	90 (+0)	727	24	90 (+0)	724	12	91 (+1)	10	4	91 (+1)	0	178 (+88)
alu-v3_35	5	19 + 18 = 37	64	103	64 (+0)	101	17	64 (+0)	74	9	64 (+0)	3	3	68 (+4)	0	121 (+57)
alu-v4_37	5	19 + 18 = 37	64	119	64 (+0)	121	17	64 (+0)	43	9	64 (+0)	6	3	68 (+4)	0	110 (+46)
mod5d1_63	5	9 + 13 = 22	48	14	48 (+0)	13	11	48 (+0)	8	7	48 (+0)	5	5	48 (+0)	1	98 (+50)
mod5mils_65	5	19 + 16 = 35	64	96	64 (+0)	98	16	64 (+0)	94	8	65 (+1)	1	3	65 (+1)	0	108 (+44)
qc_qft_4	5	44 + 27 = 71	94	136	94 (+0)	135	19	94 (+0)	21	14	94 (+0)	9	16	94 (+0)	12	115 (+21)
qc_qft_5	5	69 + 38 = 107	135	401	135 (+0)	395	26	135 (+0)	21	19	139 (+4)	107	24	145 (+10)	48	163 (+28)

n: number of logical qubits original cost: number of single qubit gates plus number of CNOT gates before mapping c: cost (number of operations) of the mapped circuit
 Δ_{min} : difference to minimum cost t: runtime in seconds

minimality. Limiting the number of permutation also has a tremendous effect on the runtime. In fact, the runtime required to solve an instance indirectly correlates with $|G'|$. However, limiting the number of permutations too much generates rather poor results regarding minimality. For the benchmarks we considered in our evaluation, the strategy *disjoint qubits* always generates results with minimum cost, whereas the strategy *qubit triangle* yields the poorest results regarding minimality. Still, all strategies provide alternatives delivering solutions that are close to the minimum within acceptable runtime.

In a second series of evaluations, we compared the obtained minimal and close-to-minimal solutions to the mapping algorithm provided by IBM’s Qiskit [12]. This allows to evaluate how far existing heuristic approaches are from the optimum. To this end, we utilized the mapper available in Qiskit 0.4.15 and list the number of gates in the obtained circuit in the last column of Table 1 (we do not list the runtime since all mappings could be determined within a second).⁸ Since the mapping algorithm in Qiskit is probabilistic, we ran it 5 times for each benchmark and list the observed minimum.

As can be seen, the heuristic approach utilized in IBM’s SDK Qiskit can be improved significantly. Considering the benchmarks *alu-v3_35* and *mod5d1_63*, IBM’s algorithm is 89% and 104% above the practical minimum that can be reached, respectively. On average, IBM’s solution yields circuits that are 45% above the minimum (by means of gate count). Considering only the number of gates added during the mapping (and not the complete mapped circuit), Qiskit’s solutions are 104% above the minimum given by \mathcal{F} on average—doubling the overhead required for mapping a circuit. Hence, even though the exact approach proposed in this paper is only applicable for mapping small quantum circuits, it shows that there is much room for improvement of heuristic approaches—further motivating research on this topic.

6 CONCLUSIONS

In this paper, we proposed an exact solution for the mapping of quantum circuits to IBM QX architectures—an \mathcal{NP} -complete problem. To this end, we formulated the considered problem symbolically and used the Boolean satisfiability solver Z3 to determine a minimal solution. Moreover, we have shown how to improve the performance by restricting the search space while still guaranteeing close-to-minimal solutions. By this, we do not only provide a method that maps quantum circuits to IBM’s QX architectures with

a *minimal* number of SWAP and H operations, but also show by experimental evaluation that the number of operations added by IBM’s heuristic solution exceeds the lower bound by more than 100% on average. An implementation of the proposed methodology is publicly available at http://iic.jku.at/eda/research/ibm_qx_mapping.

ACKNOWLEDGMENTS

This work has partially been supported by the LIT Secure and Correct System Lab funded by the State of Upper Austria and the European Union through the COST Action IC1405.

REFERENCES

- [1] M. Amy, D. Maslov, M. Mosca, and M. Roetteler. A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits. *IEEE Trans. on CAD of Integrated Circuits and Systems*, 32(6):818–830, 2013.
- [2] A. Botea, A. Kishimoto, and R. Marinescu. On the complexity of quantum circuit compilation. In *Symposium on Combinatorial Search*, 2018.
- [3] R. Courtland. Google aims for quantum computing supremacy. *IEEE Spectrum*, 54(6):9–10, 2017.
- [4] A. W. Cross, L. S. Bishop, J. A. Smolin, and J. M. Gambetta. Open quantum assembly language. *arXiv preprint arXiv:1707.03429*, 2017.
- [5] L. De Moura and N. Björner. Z3: An efficient SMT solver. In *International conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 337–340. Springer, 2008.
- [6] L. Gomes. Quantum computing: Both here and not here. *IEEE Spectrum April 2018*, 2018.
- [7] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Theory of computing*, pages 212–219, 1996.
- [8] J. Hsu. CES 2018: Intel’s 49-qubit chip shoots for quantum supremacy. *IEEE Spectrum Tech Talk*, 2018.
- [9] IBM Q team. IBM Q. <https://www.research.ibm.com/ibm-q/>. Accessed: 2018-11-26.
- [10] IBM Q team. IBM Q 5 Tenerife backend specification v1.3.0. <https://ibm.biz/qiskit-tenerife>. Accessed: 2018-11-26.
- [11] IBM Q team. Qiskit Developer Challenge. <https://qx-awards.mybluemix.net/#qiskitDeveloperChallengeAward>. Accessed: 2018-11-26.
- [12] IBM Q team. Qiskit Python SDK. <https://github.com/Qiskit/qiskit-sdk-py>. Accessed: 2018-11-26.
- [13] G. Li, Y. Ding, and Y. Xie. Tackling the qubit mapping problem for NISQ-era quantum devices. *arXiv preprint arXiv:1809.02573*, 2018.
- [14] D. M. Miller, R. Wille, and Z. Sasanian. Elementary quantum gate realizations for multiple-control Toffoli gates. In *Int’l Symp. on Multi-Valued Logic*, pages 288–293, 2011.
- [15] M. Nielsen and I. Chuang. *Quantum Computation and Quantum Information*. Cambridge Univ. Press, 2000.
- [16] J. Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, 2018.
- [17] P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. *Foundations of Computer Science*, pages 124–134, 1994.
- [18] M. Siraichi, V. F. Dos Santos, S. Collange, and F. M. Q. Pereira. Qubit allocation. In *Int’l Symp. on Code Generation and Optimization (CGO)*, pages 1–12, 2018.
- [19] R. Wille, G. Fey, D. Große, S. Eggersgluß, and R. Drechsler. SWORD: A SAT like prover using word level information. In *VLSI of System-on-Chip*, pages 88–93, 2007.
- [20] R. Wille, D. Große, L. Teuber, G. W. Dueck, and R. Drechsler. RevLib: an online resource for reversible functions and reversible circuits. In *Int’l Symp. on Multi-Valued Logic*, pages 220–225, 2008. RevLib is available at <http://www.revlib.org>.
- [21] R. Wille, A. Lye, and R. Drechsler. Exact reordering of circuit lines for nearest neighbor quantum architectures. *IEEE Trans. on CAD of Integrated Circuits and Systems*, 33(12):1818–1831, 2014.
- [22] A. Zulehner, A. Paler, and R. Wille. An efficient methodology for mapping quantum circuits to the IBM QX architectures. *IEEE Trans. on CAD of Integrated Circuits and Systems*, 2018.
- [23] A. Zulehner and R. Wille. Compiling SU(4) quantum circuits to IBM QX architectures. In *Asia and South Pacific Design Automation Conf.*, pages 185–190. ACM, 2019.

⁸Note that, to ensure a fair comparison, we only considered the actual mapping process of Qiskit and not the decomposition as well as pre- or post-mapping optimizations that may be applied before/after the mapping.