PAPER • OPEN ACCESS

NISQ circuit compilation is the travelling salesman problem on a torus

To cite this article: Alexandru Paler *et al* 2021 *Quantum Sci. Technol.* **6** 025016

View the [article online](#) for updates and enhancements.

IOP | ebooks™

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

Quantum Science and Technology

PAPER

OPEN ACCESS

RECEIVED
22 September 2020

REVISED
2 February 2021

ACCEPTED FOR PUBLICATION
15 February 2021

PUBLISHED
8 March 2021

Original content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the
title of the work, journal
citation and DOI.

NISQ circuit compilation is the travelling salesman problem on a torus

Alexandru Paler^{1,2,*} , Alwin Zulehner¹ and Robert Wille¹

¹ Institute of Integrated Circuits, Johannes, Austria

² Transilvania University of Braşov, Romania

* Author to whom any correspondence should be addressed.

E-mail: alexandrupaler@gmail.com

Keywords: NISQ, compilation, complexity

Abstract

Noisy, intermediate-scale quantum (NISQ) computers are expected to execute quantum circuits of up to a few hundred qubits. The circuits have to conform to NISQ architectural constraints regarding qubit allocation and the execution of multi-qubit gates. Quantum circuit compilation (QCC) takes a nonconforming circuit and outputs a compatible circuit. Can classical optimisation methods be used for QCC? Compilation is a known combinatorial problem shown to be solvable by two types of operations: (1) qubit allocation, and (2) gate scheduling. We show informally that the two operations form a discrete ring. The search landscape of QCC is a two dimensional discrete torus where vertices represent configurations of how circuit qubits are allocated to NISQ registers. Torus edges are weighted by the cost of scheduling circuit gates. The novelty of our approach uses the fact that a circuit's gate list is circular: compilation can start from any gate as long as all the gates will be processed, and the compiled circuit has the correct gate order. Our work bridges a theoretical and practical gap between classical circuit design automation and the emerging field of quantum circuit optimisation.

1. Introduction

The first general purpose quantum computers, which are called noisy, intermediate-scale quantum (NISQ) computers [21], operate on a few hundred qubits and do not support computational fault-tolerance. The IBM Q experience computers, which fall into the NISQ category, have sparked the interest in the automated compilation of arbitrary quantum circuits. Near-term applications of NISQ may be used to explore many-particle quantum systems or optimisation problems, and the executed circuits are not expected to include sequences longer than 100 gates [21]. Although this is a serious limitation, it is hoped that hardware quality will increase such that longer circuits may be executed.

NISQ compilation is motivated in part by the different architectures, but more important by the technical limitations of NISQ hardware, such as qubit and quantum gate fault rates, gate execution time etc. Before executing a quantum computation the corresponding circuit has to be adapted for the particularities of the NISQ computer.

1.1. Background

In order to show that the problem of quantum circuit compilation (QCC) is equivalent to a travelling salesman on a torus (e.g. figure 1) the following background material will be introduced.

We will use NISQ, quantum computer, and chip interchangeably. For the purpose of this work, a chip is described entirely by the set of hardware qubits, also called registers [28], and the set of supported interactions. The computer is abstracted by a *coupling graph* (e.g. figure 2(b)), where the registers are the vertices, and the edges are the supported multiqubit gates between vertex tuples. In a directed coupling graph $G = (V, E)$, having $|V| = q$ and $|E| \leq q(q-1)$, the edges stand for the CNOTs supported between pairs of physical qubits. The edge directions indicate which qubit is control or target. If the computer

Figure 1. The search landscape of QCC is a discrete two dimensional torus. The blue circle represents the timeline of the quantum circuit being compiled. The green vertices (only a few are drawn) are qubit allocation configurations encountered during compilation. The red circle is connecting all the qubit allocations configurations available while compiling one of the circuit's gate. QCC compilation forms a loop formed of green vertices. The solution loop intersects non-trivially all possible red loops on the torus.

Figure 2. (a) Quantum circuit example; (b) example of a coupling graph where the vertices $Q3$ and $Q4$ are not adjacent to $Q0$. Considering the graph from (b), the CNOTs $Q0 \rightarrow Q3$ and $Q0 \rightarrow Q4$ are remote. The remote CNOTs are highlighted in (a).

supports both CNOT directions between a pair of qubits, there are two directed edges between the corresponding graph vertex pairs. Current NISQ devices do not restrict CNOT direction, and G graphs are nowadays mostly undirected.

In general, NISQs do not have all-to-all connectivity between the registers, and do not support the arbitrary application of multiqubit gates. Consequently, not all the CNOTs of a circuit can be executed without further adjustment.

The QCC problem is: *for a given coupling graph and a quantum circuit C , compile a circuit C' which is functionally equivalent to C and compatible with the coupling graph.*

We use the following operations to solve QCC: (1) qubit allocation; (2) CNOT gate scheduling, and (3) circuit traversal—choosing the order in which the CNOTs are compiled. The first two are practically already methodological parts of established quantum circuit design frameworks such as Cirq and Qiskit. A theoretical analysis of the first two was provided in [27]. The third operation is based on circular CNOT circuits as introduced in [19].

Each of the three operations can be attached to an optimisation problem. Each of those problems is directly connected to the execution of a **remote CNOT**, which is defined as the gate that has to be executed between circuit qubits allocated on non-adjacent NISQ registers (e.g. figure 2). The compilation of a remote CNOT introduces additional gates into C' , because the qubits have to be effectively moved across the chip until these are on adjacent registers. A correctly compiled C' has no remote CNOTs. It is assumed that the NISQ chip has at least as many registers as C' .

We define the qubit allocation problem with respect to the effect of compiling remote CNOTs. Figure 3 includes two allocation configurations. The labels inside the coupling graph vertices represent the allocated circuit qubits. After moving Q_4 from one register to another, the configuration from figure 3(a) changes to the one from figure 3(c).

Problem 1 Qubit allocation: assign circuit qubits to NISQ registers, such that the compiled C has a minimal cost.

The cost mentioned in the problem 1 could be gate count, circuit depth etc. For example, the cost can be expressed in terms of physical CNOT gates and assuming a linear nearest neighbour architecture, in figure 2 the cost of implementing the first CNOT is zero, the second remote CNOT has a cost of six because two SWAP gates may be necessary etc.

For the purpose of this work, the compilation of remote CNOTs to the NISQ chip is a kind of gate scheduling (see problem 2) and we present an example in figure 3. Automatic approaches for gate scheduling range from global reordering of quantum wires [36] to application of circuit rewrite rules [25]. Gate scheduling has been performed even manually, by designing circuits that conform to the architectural constraints [1, 7].

Problem 2 Gate scheduling: choose the coupling graph edge where to execute a remote CNOT.

Gate scheduling is a sub-problem of the qubit allocation problem, because scheduling is performed once qubits are allocated. However, for an exact solution, finding the best initial allocation requires iterating through all possibilities which in turns implies that gate scheduling has to be calculated each time. From this perspective, QCC is at least as complex as the qubit allocation problem.

Regarding scheduling, it is not obligatory to start compiling from the first remote CNOT of C . One can start from an arbitrary gate, as long as the resulting C' will respect the original order from C . For example, if C consists of three remote CNOTs g_1 , g_2 and g_3 , the compilation could start with g_2 , followed by g_3 and finally g_1 . However, C' would need to execute the correct order of compiled gates g'_1, g'_2, g'_3 .

Problem 3 Circuit traversal: determine the order in which the gates of C should be compiled, such that the cost of C' is minimised. The chosen order has to be a valid topological sorting of C .

1.2. Related work

Most NISQ devices have a topology which is not compatible with the quantum circuits that have to be executed on them. Those circuits need to be accordingly modified. Originally, this has been done by adapting the quantum circuit in a systematic manner (e.g. [7]). However, such an approach obviously is not feasible—particularly with increasing size of the considered quantum circuit.

In the past, a huge variety of methods addressing this problem have been proposed. While some of them (e.g. [11, 23, 34, 36]) aim to solve QCC in an exact fashion (i.e. generating minimal solutions) most of them provide heuristics. Heuristics are much more established solutions for QCC, while exact approaches are mainly used for evaluation purposes (i.e. checking how far heuristics are from the optimum) or to generate quantum circuits for certain ‘building block’-functionality.

Most of the available heuristics employ a *swapping*-scheme, i.e. they insert remote CNOT and SWAP operations into the originally given quantum circuit that exchange the state of two physical qubits whenever they do not satisfy a connectivity constraint. By this, the mapping of the logical qubits of the quantum circuit to the physical ones of the hardware changes dynamically, i.e., the logical qubits are moved around on the physical ones. Approaches following this scheme include e.g. [8, 11, 15, 16, 25, 35, 37].

Other approaches use a *bridging*-scheme which does not dynamically change the mapping of the logical qubits to the physical ones: CNOT gates that violate the connection constraint are decomposed into several CNOT gates that bridge the ‘gap’. This scheme has the advantage that, given the initial mapping, determining the mapped circuit is straightforward. On the other side, it often leads to more costly solutions since the number of CNOT operations required to realize bridge gates grows exponentially. Approaches following this scheme include e.g. [4, 5, 10, 11, 23].

1.3. Complexity of QCC

Multiple approaches to showing the complexity of QCC have been presented. One of the first, Maslov [14] demonstrated that a variant of QCC is NP-complete by showing that it implies the search of a Hamiltonian cycle in a graph. In the context of our QCC formulation, the work of [14] is concerned with optimal solutions to the qubit allocation problem when the blue torus edge weights are determined by the physical gate execution times along the longest input–output gate chain.

QCC has also been considered a search problem according to [27], which includes a detailed review of the methods used for determining the complexity class. It has been recently discussed that the complexity of QCC optimisation is NP-hard [17] by comparing QCC with the optimisation of fault-tolerant quantum circuits protected by the surface code [9].

The authors of [2] have shown that QCC as a discrete optimisation of a circuit's makespan is NP-complete for QAOA circuits. The proof from [2] is a reduction from the Boolean satisfiability problem (SAT) to the QCC problem, and was applied to circuits consisting of two types of two-qubit gates: SWAP and PS (phase separation). The proof did not rely on any particular ordering of the gates in the circuit. Such circuits can be decomposed with a constant overhead into the circuits we consider in this work (CNOT gates and single qubit gates). PS gates can be decomposed using the KAK decomposition [32, 38] into CNOTs and single qubit gates, and the SWAP gate can be decomposed into three CNOTs.

Moreover, a method for optimising QAOA circuits by taking commutativity into account was presented by [31]. Therein the authors show very convincingly that theorem proving (e.g. Z3 solver) and SAT solvers do not scale for practically large compilation problems (more than 100 qubits and deep circuits): the search space of QCC as an NP-complete problem is still exponential even when the number of variables is reduced exponentially. The work of [31] combined with the theoretical approach from [2] highlight the importance of QCC heuristics.

QCC has been presented as an application of temporal planning [33], too. In general, temporal planning can have a higher complexity than NP-complete. For example, concurrent temporal planning is EXPSPACE-complete [24]. This, however, does not imply that QCC would be EXPSPACE-complete. In fact, in domain-independent planning, it is not uncommon that a planning system attacks a domain with a lower complexity than the complexity of the AI planning variant that the planning system at hand can handle. This is done for the convenience of using a readily available off-the-shelf system, when a domain-specific solver is not necessarily available.

2. Methods

We present the construction of how QCC can be solved as the travelling salesman problem (TSP). To this end, we illustrate the construction of the QCC torus.

2.1. Arranging qubit allocations in a circle

Allocating circuit qubits to NISQ registers can be expressed as a permutation vector of length q . For example, assume that Q_i are the qubits of a circuit C , and H_i are the registers of a computer, for $i \leq q$. Both the computer and the circuit have $q = 5$ qubits. The permutation $p_1 = (0, 1, 2, 3, 4)$ is the trivial allocation where Q_i are allocated to $p_1[Q_i] = H_i$: circuit qubit Q_0 at register H_0 , qubit Q_1 at register H_1 etc. Another example is the permutation $p_2 = (2, 1, 0, 4, 3)$, where Q_0 is allocated to H_2 , Q_1 is allocated to H_1 etc.

In the following, a **configuration** is a permutation that represents how circuit qubits are allocated to NISQ registers. The terms permutation and configuration will be used interchangeably. For example, p_1 and p_2 are configurations, too.

The set of all permutations forms a symmetric group with $q!$ elements. The group has $q - 1$ transposition generators. A transposition swaps two elements of the permutation, and keeps all other entries unchanged. Any group element is expressed through a non-unique sequence of transposition generators.

The group structure can be visualised as a graph. The elements are vertices, and edges are transpositions connecting the vertices. If all group elements are exhaustively enumerated, the graph is a circle (e.g. figure 4) with $q!$ vertices. There exist more compact representations of the group, such as the complete graph K_q . Without affecting the generality, the exhaustive representation is preferred in this work.

2.2. The circuit as a circle of CNOTs

The compilation problem has been reduced to scheduling the execution of CNOTs, remote or not. Quantum circuits are often manipulated as directed acyclic graphs (DAGs) with vertices for quantum gates. Edge directions reflect the gate ordering inside the circuit. For the purpose of this work, the DAG representation is replaced by the equivalent (blue) circle of CNOTs [19]. The order of the vertices on the

Figure 4. Circular arrangement of $6 = 3!$ configurations. In order to avoid confusion with a coupling graph, vertices are polygons. The permutation vector associated to each vertex is placed next to the vertex. For example, $[0, 1, 2]$ is next to p_0 . The edge represents transpositions. For example, the permutation p_0 is transformed into p_1 by transposing 0 and 1. The edges and vertices are red in order to highlight the correspondence to figure 1. The weights along the red circles are zero (section 2.4).

Figure 5. Example: a quantum circuit with four CNOTs named $G_0, \dots, G_3$ is drawn as a circular CNOT circuit having the wires as concentric circles. The normal circuit is used to obtain a chain of configurations after the application of each CNOT. The edges in the chain represent the CNOT. The ends of the chain can be joined to form a circle. The edges and vertices are blue in order to highlight the correspondence to figure 1.

blue circle encodes one of the equivalent topological orderings from the DAG. In general, gate commutativity may be used to improve the compiled circuit (see appendix A on the backtracking method). In particular, all equivalent DAG topological orderings may need to be considered. The latter is equivalent to commuting gates from the chosen topological ordering with an identity gate.

A circle is obtained as follows: (a) only the CNOTs are kept from the circuit, and other gates are discarded (e.g. the T gates from figure 4), (b) the wire endpoints corresponding to input and output are joined together. Figure 5 is an example of obtaining a circular CNOT circuit. Pairs of adjacent vertices in the chain represent the qubit allocation configurations before and after a remote CNOT was compiled.

It is possible to start compiling a circuit C from any gate g and not necessarily from the first gate. The circular CNOT circuit supports the correctness of this observation. Let us consider that the circuit C is the application of a sequence of two sub-circuits A and B , such that $C = AB$. Moreover, we model QCC as a function that computes $QCC(C) = C'$, where $QCC(C) = QCC(A)QCC(B) = C'$.

Instead of starting with the first gate of A , we assume that compilation starts from sub-circuit B and that the CNOT circle is traversed in the correct order. Along the circular traversal $A^\dagger$ will be compiled instead of A . The compilation result will be circuit D for which $D = QCC(B)QCC(A^\dagger)$.

However, it is possible to reconstruct C' by inverting the gate list of $QCC(A^\dagger)$ such that $QCC(A^\dagger)^\dagger = QCC(A)$. This divide and conquer approach does not imply that a greedy approach can solve QCC efficiently. It still is combinatorially difficult to choose the *best* gate from C to start compilation from.

Starting the traversal of CNOT circles from arbitrary positions can be advantageous for reducing the total cost of the compiled circuit.

It is not guaranteed that $\text{Cost}(C, D) = \text{Cost}(C, C')$, such that a heuristic approach to QCC could be to start compiling from different gates of C .

Figure 6. After cutting the torus along the yellow line, the start circle is brown and the stop circle is red. Compilation is the problem of connecting the circles by a loop of green vertices, such that the sum of the edge costs is minimal.

Figure 7. Example of a search diagram with $3!$ configurations. The red configuration circles are concentric. The red circles are interconnected by CNOT circles that were cut like in figure 6. Search starts in the centre of the diagram and stops at one of the endpoints of the radial CNOT circles. The green vertices indicate the path found by the compiler.

2.3. Unfolding the torus

The circular graph of configurations and the CNOT circle can be combined to a *torus* (e.g. figures 1 and 6). The torus has a discrete structure, which can be used to visualise and analyse QCC. For visualisation purposes, the torus can be cut and unfolded to a planar structure. We will resort to a single cut along the configuration circle. The result will be a two dimensional diagram like the one in figure 7. Let one side of the cut be called the start-circle and the other side the stop-circle.

As shown in figures 7 and 8, a hypothetical quantum compiler will traverse vertices of the torus. The number of torus vertices is the total number of states the compiler should consider, and there are $q! \times |C|$ states. By restating problem 1, the compiler will find a path from the start circle to the stop circle (figures 7 and 6). There is a combinatorial number of paths of various lengths between pairs of start-stop vertices. QCC executes, in the best case, linearly in the number of circles traversed between start and stop.

2.4. Edge weights

The edges connecting the vertices of the torus are weighted. Two extreme cases are possible: (a) all edges have weight zero; (b) all edges have equal weight. The first case is not realistic in the context of QCC. The

second case arises when the NISQ device has all-to-all connectivity, such that the shortest path between a start and a stop circle is given by the straight traversal of a CNOT circle. For the purpose of this work, the red edges (configuration edges) have zero weight, and the blue edges (CNOT edges) have non-zero weight. The motivation for this model is twofold.

First, our goal is to show that QCC is TSP (see section 3.1), and we have chosen the generalised TSP (GTSP) form as presented in [18]. This TSP form uses the concept of connected city clusters. The movement within a cluster has cost zero, but the movement between clusters non-zero. In our case, the red rings are the clusters and the blue rings are the connections between the clusters.

Second, instead of weighing the red edges, we consider their cost as part of the blue edge weights. Each blue edge traversal requires compiling a (remote) CNOT. The compilation is thus determined by: (a) the cost of implementing the transposition resulting by moving along the red ring (a new start qubit allocation configuration from which the CNOT is compiled), and (b) the cost of effectively scheduling the remote CNOT.

Additionally, we note that by joining the first and last red rings the qubit allocation configuration has to be the same, in general. This is the case, when compilation does not start from the first gate of the circuit (cf section 2.2) and needs to reconstruct the solution. After reconstructing the solution, however, the wire permutations before the first and last gate can be removed—these are simple wire relabelling operations. As a result, configuration changes on the start/stop circle come for free and are not considered in the compilation cost, because no gates need to be inserted in the circuit. For example, in figure 8 the orange traversal of the configuration ring has cost zero.

This brings us to the particular QCC scenario, which we assume being the common one, when compilation starts from the start ring (e.g. brown in figure 6 corresponds to the first gate from the uncompiled circuit) and ends on a different vertex of the same ring. Different vertices on the same ring refer to different qubit allocation configurations. Therefore, in particular, it is an acceptable solution to end on the same ring, but on a different vertex.

We mentioned that the weights may be, for example, the number of physical CNOTs necessary to implement a remote CNOT. In general, edge weights are assigned by a **cost function**. It is the task of the cost function, for example, to perform topological analysis of the circuit and coupling graph [6]. The cost of gate compilation could include also lookahead information, similarly to how this was performed for example for linear nearest neighbour architectures [35].

Formulating explicit cost functions does not fall within the scope of this work. As shown in [31], once the cost functions are specified, formulating the optimisation objective is a highly nontrivial task. The optimisation objective is for exact QCC methods like the actual code implementation is to the heuristic QCC methods. Therefore, even if we would specify the exact functions, the optimality of the compiled circuit would depend on the time-space trade off allowed by the heuristic implementation. In particular, just as examples: (a) if the optimisation goal is the minimum number of SWAPs one could use the MI strategy from the appendix A; (b) for minimising depth, and by making no assumptions about gate execution time like in [14], the optimisation goal would be makespan [2].

From the perspective of an arbitrary function $\text{Cost}(C, C')$ that calculates the cost of compiling C into C' (similar to discussion in (section 3.1)), we can state that QCC optimisation is to find a circuit M such that $\text{Cost}(C, M) = \min(\text{Cost}(C, C_i))$, where C_i is a loop on the torus. The best C_i loop has the minimum sum of the traversed edges.

3. Results

The landscape of QCC is a discrete torus obtained from the Cartesian product of two circles. One of the circles refers to the group structure of the qubit allocations possible when scheduling a gate (i.e. the red circle in figure 1). The other circle is generated by the fact that the CNOTs of a circuit can be arranged in a circular form (i.e. the blue circle in figure 1) [19].

The torus includes $|C|$ red circles—one for each gate from C . There are $q!$ blue circles: one for each possible permutation of circuit qubits to NISQ registers. The details of constructing the torus were presented in section 2.

3.1. QCC is a TSP

In the following we show that QCC is a TSP. The appendix A includes a backtracking formulation of QCC as TSP. Independent of this work, the authors of [27] have decomposed the compilation problem into two steps: qubit allocation and scheduling of multi-qubit gates. In practice, this approach has already been followed by quantum circuit frameworks such as Cirq and Qiskit: the circuit qubits are mapped to the NISQ device, and then the circuit gates are scheduled. For QCC benchmarking purposes, the two-step approach has also been used by [30]. We augment the QCC decomposition by including the circular CNOT structure. This will be useful for analysing the problem complexity. The exact complexity depends on how the cost function is implemented and evaluated.

We use the following definitions:

- **A solution** is any loop that intersects non-trivially the red circle from figure 1. There is a combinatorial number of potential solutions.
- **The minimum solution** is the loop for which the sum of the edge weights is minimal (section 2.4).

According to the discussion in section 2.4, only the edges along the CNOT circles have non-zero weights. Each solution is the sum of $|C|$ weights $\text{Cost}_l = \sum_0^{|C|} w_{p,q}$, where l is the index of the solution and p, q are the indices of the configurations connected by the edge that has weight $w_{p,q}$. The solution of QCC is $\min(\text{Cost}_l)$, for $l \leq (q!)^{|C|}$ when the exhaustive enumeration of the configurations is used. In the light of the definitions of problems 1–3, where a minimum cost circuit is searched, QCC is an example of combinatorial optimisation.

We show that QCC is a GTSP. The original TSP problem is defined for a number of cities, for which the distances between pairwise cities are known. TSP answers the question: what is the shortest possible route visiting all the cities and returning to the origin city? In GTSP the cities are arranged into clusters, and the edges connecting the cities inside the cluster have weight zero [18]. At least one city from each cluster has to be visited on the shortest path [18].

QCC is GTSP when considering each red configuration ring of the $|C|$ as a cluster of cities. Moreover, the zero weight cluster edges are consistent to how the weights along the configuration rings are set in section 2.4. There are $|C|$ configurations circles in the torus. The distances between the cities are the weights along the CNOT edges. The salesman is expected to traverse at least once each configuration circle between the red start circle and the brown circle from figure 6.

The fact that the configuration rings are arranged in a circle does not make the problem easier. Assuming that $|C|$ has only three remote CNOTs, then there are only three clusters for which the GTSP has to be computed. However, the arrangement of the three clusters corresponds to a complete graph K_3 —the smallest instance of GTSP. Increasing the length of the circuit increases the number of clusters, but does not reduce the complexity of the optimisation problem.

The decision GTSP version of QCC answers the question: is there a route/loop of cost less than a specified $\text{Cost}_{\text{route}}$? Any potential solution can be verified by tracking the proposed solution loop along the torus. Because of its complexity, QCC has to be solved using heuristics. Benchmarking circuits for which the minimum $\text{Cost}_{\text{route}}$ is known beforehand [30] are a good way to evaluate the performance of the heuristics.

3.2. QCC is a ring

The discrete torus shows that QCC, from the perspective of discrete mathematics, is a ring with the two QCC-operations being: (1) qubit allocation; (2) gate scheduling.

We can define commutativity in a manner compatible with quantum circuit execution. Very informally, two QCC-operations are commutative iff the computation implemented by the circuit is unchanged after reordering the QCC-operations. Consequently, qubit allocation is commutative because it is effectively a renaming of wires. It does not matter in which order the qubits are allocated, this does not change the computation. QCC-gate scheduling is not commutative because, in general, two CNOTs are not commutative. Consequently, the circle of allocations is the illustration of an Abelian group, and the CNOT-circle represents a monoid. The Abelian group and the monoid form the discrete torus where traversal are unidirectional. We leave a formalisation of the mathematical structure of QCC for future work.

4. Discussion and conclusion

NISQ compilation is receiving increased attention, due to its practical industrial relevance. In this work, the QCC problem was decomposed into a set of sub-problems, whose individual solution is found by traversing circles. This enabled the formulation of QCC as a TSP along a torus. We have implemented the TSP approach to QCC compilation at <https://github.com/alexandrupaler/k7m> and we have used it as part of a machine learning approach to QCC in [20].

The torus structure presented has the potential to generate other efficient heuristics for the QCC compilation. Exact QCC methods [31, 36] scale poorly, because these are as fast as the underlying solver. The highly regular and cyclic structure of the torus search space may inspire improved variable encodings such that exact layout methods can be pushed in the area of 100-qubit circuits.

The TSP formulation hints at the conceptual similarities between QCC and the automatic design of quantum optical experiments [12]. The latter consist of discrete optical elements, which can be placed in a combinatorial (mapping steps) number of experimental configurations formed by different devices (scheduling step). At the same time, forming loops on the torus shows that QCC is similar to a dynamic optimisation problem [3], and that it would be reasonable to expect methods based on ant colonies or evolutionary algorithms for solving QCC.

Finally, because quantum methods have been applied to TSP (using QAOA in [22], using Grover's algorithm by [29]), our work opens the possibility to optimise quantum circuits with quantum computers.

Acknowledgments

We are grateful to Adi Botea for feedback and suggestions to improve the manuscript. We acknowledge the input of Razvan Andonie on optimisation problem complexity classes, as well as the corrections proposed by Daniel Herr. This work was funded by the Linz Institute of Technology project CHARON, the Google Faculty Research Award FRA_ANGELICO, and the NUQAT project of Universitatea Transilvania Brasov.

Data availability statement

No new data were created or analysed in this study.

Appendix A

A.1. Moving on a (red) configurations circle

At least two strategies are possible for compiling a *single remote CNOT*. The methods are illustrated in figure 9. The first strategy is MIM (abrv. for move-interact-move): move one of the qubit states on a wire next to the other qubit's wire, interact the qubits, and then swap back to the original wire. The second strategy is called MI (abrv. for move-interact) and is similar to the first one but without swapping back the moved qubit state.

Applying MIM once introduces $2d$ SWAP gates in the circuit, while the MI strategy only d SWAPS, where d is the distance between the remote wires. A straightforward distance function could be, for example, the Manhattan distance which can be used for LNN as well as grid NISQ architectures. For a given permutation p , after applying MIM, the resulting permutation is also p .

On the contrary, after an MI swap, the resulting permutation is a p' , obtained through the sequence of transpositions representing the SWAP gates. Although MI introduces less SWAPS, it increases the complexity of the compilation problem: each remote CNOT will result in a new permutation, such that the circuit qubit allocation configuration is *evolving* after each CNOT.

In the presence of evolving configurations, state of the art compilation methods are solving the following problem: find an optimal circuit consisting entirely of SWAP gates that transforms a current permutation p_{in} to a permutation p_{out} such that a given batch of remote CNOTs can be implemented on the given coupling graph. In other words, an optimal sequence of transpositions is sought, such that p_{out} conforms to a set of constraints imposed by all the CNOTs to implement. During the search of a SWAP circuit, or after a SWAP circuit was found, it is checked that p_{out} conforms to the coupling graph.

This approach implies that p_{out} and the SWAP circuit generating it are computed for more than a single CNOT (e.g. figure 10). Both the set of remote CNOTs and the SWAP circuits are computed using heuristics (e.g. randomised algorithm in the IBM QISKit, A*-search [37] or temporal planners [33]).

A.2. Moving on a (blue) CNOT circle

Movement on a CNOT circle is equivalent to compiling CNOT gates sequentially. This is not to say that CNOT cannot be parallelised in the resulting C' circuit. Parallelisation of a batch of CNOTs can be visualised on the torus: a CNOT is selected from the batch and compiled such that, for the remaining CNOTS, zero weights are placed on the edges connecting the configuration circles. Thus, for the first CNOT a kind of lookahead strategy [35] has to be used to determine the configuration that will generate zero weight edges in the future.

Without discussing lookahead methods, compilation implies finding a *good* configuration and then advancing on the CNOT circle. Thus, compilation is preceded by movements along the configuration circle whenever SWAP networks are used to prepare the configuration. But because remote CNOTs can be implemented also without SWAP networks, compilation of remote CNOTs can also have a different cost.

A.3. Backtracking for TSP

Having paralleled QCC to TSP, we can formulate a naive backtracking algorithm for compilation. The first step of the algorithm is to determine an initial configuration: how circuit qubits are mapped (allocated) to the NISQ (figure 8(a)). Afterwards, the first edge of the CNOT circle starting from this configuration vertex is traversed by choosing a coupling graph edge where to execute the CNOT. A new configuration is reached by using the MI swap strategy (figure 11). The next torus edge traversal is prepared by moving around the configurations circle (figure 8(b)) and landing in a new configuration.

The backtracking step consists of two sub-steps: traversing the current configuration circle, followed by traversing the CNOT circle. The backtracking step undoes the last CNOT compilation and moves along the previous configuration circle. This is equivalent to selecting a different edge where to map the remote CNOT that was just undone.

A solution is found each time a vertex from the outmost CNOT circle, marked by *Stop*, is touched. Each solution is stored, and the best one is selected after the backtracking algorithm finishes: when all the cycles and configurations were naively considered.

Similarly to [31], it is possible to further increase to generality of the backtracking procedure by considering gate commutations on the blue rings figure 2 and 12. Then for each combination of the supported gate commutations, the torus has to be regenerated and the QCC procedure will have to be repeated.

A.4. Pre- and post-processing

The problem statement of QCC does not mention if C is expressed using the universal gate set supported by the NISQ. If this is not the case, C has to be translated to a functionally equivalent C' that uses gates compatible with the NISQ gate set. This is a complex QCC pre-processing task with regard to the optimal number of resulting gates (e.g. [26]), and does not fall within the scope of this work. Also, quantum algorithm and quantum hardware optimisations (cf [13]) are not considered parts of the general QCC framework.

The very high complexity of the exact method is a motivation for heuristics. It is useful to attempt to identify heuristic types and functionalities. As mentioned in section 1, compilation is the process of transforming a circuit C into another circuit C' that conforms to a set of constraints encoded into a coupling graph. Therefore, it is possible to *preprocess* C and *postprocess* C' .

Preprocessing adapts C for compilation, and it is viable to try and reduce the number of single qubit gates and CNOT gates by using, for example, template based optimisations [25]. Postprocessing can be template based too, as well as include recompilation of subcircuits of C' . For example, the IBM Qiskit uses this approach for single qubit gates, and this procedure was used by [38].

Heuristics can be included also for the previously discussed mapping problems. Selecting the start configuration (or any other configuration along the concentric cycles) could be performed using existing LNN optimisation methods, but cost models adapted to MI swaps should be formulated and analysed first. Another possibility is to collect all configurations generated along a CNOT-chain and try them out as start configurations. However, given the dimension of each configuration cycle, the collected configurations may be as good/bad as the initial one. Ranking coupling graph nodes is another heuristic for building the initial configuration [6]. The circuit mapping strategy presented in [14] would also fall in this category.

Traversal of edges along CNOT circles could be sped up by reducing the number of backtracking steps (minimum is zero), and to select from a few best possible edges for the mapping. The procedure for selecting *the best coupling graph edge* is the following: (1) shortest paths between all pairs of coupling graph vertices are computed using the Floyd–Warshall algorithm; (2) it is possible to add weights to the coupling graph edges (e.g. to prefer certain areas of the graph), or to treat the coupling graph as undirected; (3) once a remote CNOT needs to be mapped to an edge, the sum of the distances between the coupling graph vertices where the qubits are located and each graph edge vertices is computed (e.g. figure 3). The edge with the minimum distance sum is chosen, and, if multiple edges have the same distances, the last one in the list is chosen. Thus, the weighting function used for the coupling graph edges influences the edge selection).

Edge mapping could be performed for multiple remote CNOTs in parallel, too. This possibility shows that the algorithm from [37] is a heuristic fitting in the framework of this work.

Figure 12. The order of the concentric configuration circles is swapped after commuting CNOTs from C. In this example, there are two CNOTs to be scheduled: the first one is marked by green vertices, the second one by thick black stroked vertices (a single vertex of this CNOT is included in the figure). (a) The green CNOT is compiled first, and the white 1 s; (b) assuming that the CNOTs can be commuted in the original circuit, the order of the vertices in each CNOT circle can be permuted. The white CNOT is compiled first and the green CNOT second (now, a single vertex of this CNOT is included in the figure).

ORCID iDs

Alexandru Paler <https://orcid.org/0000-0002-1536-8858>

References

- [1] Boixo S, Isakov S V, Smelyanskiy V N, Ryan B, Ding N, Zhang J, Bremner M J, Martinis J M and Neven H 2018 Characterizing quantum supremacy in near-term devices *Nat. Phys.* **14** 595–600
- [2] Botea A, Kishimoto A and Marinescu R 2018 On the complexity of quantum circuit compilation *11th Annual Symp. on Combinatorial Search*
- [3] Cruz C, González J R and Pelta D A 2011 Optimization in dynamic environments: a survey on problems, methods and measures *Soft Comput.* **15** 1427–48
- [4] Alexandre A d A, Dueck G W and Alexandre C R d S 2019 CNOT gate mappings to Clifford + T circuits in IBM architectures *Int'l Symp. On Multi-Valued Logic* pp 7–12
- [5] Dueck G W, Pathak A, Rahman M M, Shukla A and Banerjee A 2018 Optimization of circuits for IBM's five-qubit quantum computers *EUROMICRO Symp. On Digital System Design* pp 680–4
- [6] Ferrari D and Amoretti M 2018 Demonstration of envariance and parity learning on the IBM 16 qubit processor *arXiv Preprint arXiv:1801.02363*
- [7] Fowler A G, Devitt S J and Hollenberg L C L 2004 Implementation of Shor's algorithm on a linear nearest neighbour qubit array *Quantum Inf. Comput.* **4** pp 237–51
- [8] Hattori W and Yamashita S 2018 Quantum circuit optimization by changing the gate order for 2D nearest neighbor architectures *Workshop on Reversible Computation* pp 228–43
- [9] Herr D, Nori F and Devitt S J 2017 Optimization of lattice surgery is NP-hard *Npj Quantum Information* **3** 1–5
- [10] Itoko T, Raymond R, Imamichi T and Matsuo A 2020 Optimization of quantum circuit mapping using gate transformation and commutation *Integration* **70** 43–50
- [11] Itoko T, Raymond R, Imamichi T, Matsuo A and Cross A W 2019 Quantum circuit compilers using gate commutation rules *Asia and South Pacific Design Automation Conf.* pp 191–6
- [12] Krenn M, Erhard M and Zeilinger A 2020 Computer-inspired quantum experiments *arXiv Preprint arXiv:2002.09970*
- [13] Maslov D 2017 Basic circuit compilation techniques for an ion-trap quantum machine *New J. Phys.* **19** 023035
- [14] Maslov D, Falconer S M and Mosca M 2008 Quantum circuit placement *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* **27** 752–63
- [15] Matsuo A, Hattori W and Yamashita S 2019 Reducing the overhead of mapping quantum circuits to IBM Q system *IEEE Int. Symp. on Circuits and Systems*
- [16] Matsuo A and Yamashita S 2019 An efficient method for quantum circuit placement problem on a 2D grid *Workshop on Reversible Computation* pp 162–8
- [17] Nash B, Gheorghiu V and Mosca M 2020 Quantum circuit optimizations for NISQ architectures *Quantum Sci. Technol.* **5** 025010
- [18] Noon C E and Bean J C 1993 An efficient transformation of the generalized traveling salesman problem *INFOR: Information Systems and Operational Research* **31** 39–44
- [19] Paler A 2016 Circular CNOT circuits: definition, analysis and application to fault-tolerant quantum circuits *Int. Conf. on Reversible Computation* (Berlin: Springer) pp 199–212
- [20] Paler A, Sasu L M, Florea A and Andonie R 2020 Machine learning optimization of quantum circuit layouts *arXiv Preprint arXiv:2007.14608*
- [21] Preskill J 2018 Quantum computing in the NISQ era and beyond *Quantum* **2** 79
- [22] Radzihovsky M, Murphy J and Mason S 2019 A QAOA Solution to the Traveling Salesman Problem Using Pyquil
- [23] Mazder Rahman M and Dueck G W 2015 Synthesis of linear nearest neighbour quantum circuits *10th International Workshop on Boolean Problems* 2012 (Springer) Freiburg, Germany
- [24] Rintanen J 2007 Complexity of concurrent temporal planning *ICAPS* **7** 280–7
- [25] Saeedi M, Wille R and Drechsler R 2011 Synthesis of quantum circuits for linear nearest neighbour architectures *Quantum Inf. Process* **10** 355–77

- [26] Shende V V and Markov I L 2009 On the CNOT-cost of TOFFOLI gates *Quantum Info. Comput.* **9** 461–86
- [27] Siraichi M Y, Santos V F d, Collange C and Pereira F M Q 2019 Qubit allocation as a combination of subgraph isomorphism and token swapping *Proc. ACM Program. Lang.* **3** 1–29
- [28] Siraichi M Y, Fernandes dos Santos V, Collange S and Quintão Pereira F M 2018 Qubit allocation *Proc. of the 2018 Int. Symp. on Code Generation and Optimization* pp 113–25
- [29] Srinivasan K, Satyajit S, Behera B K and Panigrahi P K 2018 Efficient quantum algorithm for solving travelling salesman problem: an IBM quantum experience *arXiv Preprint* arXiv:[1805.10928](https://arxiv.org/abs/1805.10928)
- [30] Tan B and Cong J 2020 Optimality study of existing quantum computing layout synthesis tools *IEEE Trans. Comput.* **1**
- [31] Tan B and Cong J 2020 Optimal layout synthesis for quantum computing *2020 IEEE/ACM Int. Conf. On Computer Aided Design (ICCAD)* (Piscataway, NJ: IEEE) p 1
- [32] Tucci R R 2005 An introduction to Cartan’s KAK decomposition for qc programmers *arXiv Preprint* arXiv:[quant-ph/0507171](https://arxiv.org/abs/quant-ph/0507171)
- [33] Venturelli D, Do M, Rieffel E and Frank J 2018 Compiling quantum circuits to realistic hardware architectures using temporal planners *Quantum Sci. Technol.* **3** 025004
- [34] Wille R, Burgholzer L and Zulehner A 2019 Mapping quantum circuits to IBM QX architectures using the minimal number of SWAP and H operations *Design Automation Conf.* p 142
- [35] Wille R, Oliver K, Walter M, Rohrs P, Chattopadhyay A and Drechsler R 2016 Look-ahead schemes for nearest neighbor optimization of 1d and 2d quantum circuits *Design Automation Conf. (ASP-DAC), 2016 21st Asia and South Pacific* (Piscataway, NJ: IEEE) pp 292–7
- [36] Wille R, Lye A and Drechsler R 2014 Exact reordering of circuit lines for nearest neighbor quantum architectures *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.* **33** 1818–31
- [37] Zulehner A, Paler A and Wille R 2018 An efficient methodology for mapping quantum circuits to the IBM QX architectures *IEEE Trans. on CAD of Integrated Circuits and Systems*
- [38] Zulehner A and Wille R 2019 Compiling SU (4) quantum circuits to IBM QX architectures *Proc. of the 24th Asia and South Pacific Design Automation Conf.* pp 185–90