

# MQT QMAP: Efficient Quantum Circuit Mapping

(Invited Summary Paper)

Robert Wille

Chair for Design Automation  
Technical University of Munich  
Munich, Germany  
Software Competence Center Hagenberg GmbH  
Hagenberg, Austria  
robert.wille@tum.de

Lukas Burgholzer

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

## ABSTRACT

Quantum computing is an emerging technology that has the potential to revolutionize fields such as cryptography, machine learning, optimization, and quantum simulation. However, a major challenge in the realization of quantum algorithms on actual machines is ensuring that the gates in a quantum circuit (i.e., corresponding operations) match the topology of a targeted architecture so that the circuit can be executed while, at the same time, the resulting costs (e.g., in terms of the number of additionally introduced gates, fidelity, etc.) are kept low. This is known as the quantum circuit mapping problem. This summary paper provides an overview of QMAP—an open-source tool that is part of the *Munich Quantum Toolkit* (MQT) and offers efficient, automated, and accessible methods for tackling this problem. To this end, the paper first briefly reviews the problem. Afterwards, it shows how QMAP can be used to efficiently map quantum circuits to quantum computing architectures from both a user’s and a developer’s perspective. QMAP is publicly available as open-source at <https://github.com/cda-tum/qmap>.

## CCS CONCEPTS

• **Hardware** → **Quantum computation; Emerging languages and compilers**; *Physical design (EDA); Hardware description languages and compilation*; • **Software and its engineering** → **Compilers**.

## KEYWORDS

quantum computing, quantum circuit mapping, compilation

## 1 INTRODUCTION

Quantum computing [1] employs quantum-mechanical principles to process data. Unlike classical computing, which uses binary digits (bits) that can only be 0 or 1, quantum computing uses quantum bits. This allows for states in superposition or states that are entangled. Because of these quantum-mechanical properties, quantum computers can solve some computational problems exponentially faster than classical computers. One of the most notable tasks is determining the prime factors of a large composite number [2]. Because many modern encryption systems are based on this problem, resolving it quickly would have far-reaching consequences for code-breaking and cybersecurity. Furthermore, quantum computing has the potential to transform fields such as machine learning [3], optimization [4], and quantum

simulation [5]. The ability to perform specific tasks exponentially faster than traditional computers may also aid in drug discovery [6], materials science [7], and artificial intelligence [8]. However, quantum error correction, the development of quantum algorithms and software, as well as the physical realization of large-scale quantum computers, continue to be major challenges for the field. But quantum computing research is moving quickly, and quantum computer capabilities are growing continuously.

Commonly, quantum algorithms are described as quantum circuits, which are sequences of quantum operations (or gates) acting on qubits. A quantum circuit compiler must be used to compile an algorithm before it can be executed on a real *Quantum Processing Unit* (QPU). Typically, such a compiler begins by simplifying the original circuit description using device-independent optimizations [9]–[14]. These optimizations include techniques like canceling or rearranging gates and using high-level synthesis routines. Following these preliminary optimizations, QPU-specific circuits are generated. This is accomplished by translating all non-native gates into a set of native gates that can be directly executed on the device [15]–[18]. Furthermore, some QPUs only feature limited connectivity between their qubits—necessitating a mapping step that maps the circuit’s logical (or virtual) qubits to the device’s (physical) qubits so that connectivity constraints are met throughout the circuit.

This work focuses on this quantum circuit mapping step. In fact, it is extremely difficult to ensure that all gates in a quantum circuit match the topology of a specified architecture while, at the same time, keeping the resulting costs (e.g., in terms of the number of additionally introduced gates, fidelity, etc.) low. In practice, most quantum circuits cannot be easily mapped to existing quantum architectures. As a result, additional gates (known as SWAP gates) are introduced during the mapping process to ensure that the circuit adheres to the architecture’s connectivity constraints. The quality of the resulting mapping can be assessed via a variety of metrics, such as the resulting circuit’s (two-qubit) gate count, depth, or expected fidelity. For the sake of clarity, we will focus on the number of additional gates in the following. It is crucial for the performance of the resulting circuit to keep this value as low as possible, as it can frequently mean the difference between a circuit producing expected results and a circuit producing entirely random outputs.

This paper shows how to tackle the quantum circuit mapping problem efficiently using the open-source tool QMAP, which is part of the *Munich Quantum Toolkit* (MQT)—a collection of

quantum computing software tools whose main focus is on providing efficient, automated, and accessible methods for tackling challenging problems in the design of quantum computing applications. To that end, we will provide a brief overview of various aspects, specifically

- (1) a high-level description of the quantum circuit mapping problem and how the MQT QMAP tool helps solve it efficiently (covered in Section 3),
- (2) the user’s perspective on how QMAP can be used to efficiently map quantum circuits to quantum computing architectures (covered in Section 4), as well as,
- (3) the developer’s perspective on how to develop new or extend existing methods in the QMAP tool (covered in Section 5).

## 2 BACKGROUND

Before delving into quantum circuit mapping and QMAP, this section provides a quick review of the fundamentals of quantum computing in general and current quantum computing architectures in particular. It should be noted that the scope of this summary paper does not allow for a comprehensive review of the broad field of quantum computation. We therefore direct anyone wishing to enter the field to a more in-depth treatment of the fundamentals as provided, for example, in [1].

### 2.1 Quantum Computing

In contrast to classical computing, which represents data using bits that can either be zero or one, quantum computing uses qubits, or quantum bits that can exist in a superposition of states, which means that they can be in a linear combination of the 0 and 1 states. More specifically, the state  $|\psi\rangle$  of a single qubit is described by two complex-valued amplitudes  $\alpha_0$  and  $\alpha_1$  such that

$$|\psi\rangle = \alpha_0|0\rangle + \alpha_1|1\rangle \text{ with } |\alpha_0|^2 + |\alpha_1|^2 = 1.$$

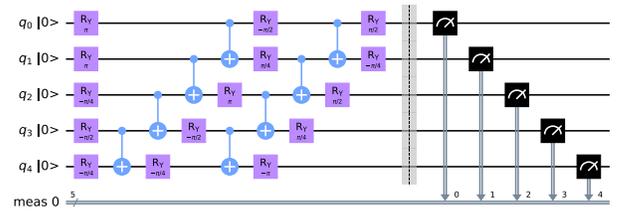
Measuring a qubit (the only way to extract information from a quantum computer), probabilistically collapses its state to one of the basis states  $|0\rangle$  or  $|1\rangle$  (with probability  $|\alpha_0|^2$  or  $|\alpha_1|^2$ , respectively).

Just as in classical computing, multiple qubits are used to perform computations. To this end, the state of an  $n$ -qubit system is described by  $2^n$  complex-valued amplitudes  $\alpha_i$ ; one for each possible bitstring of length  $n$ , i.e.,

$$|\psi\rangle = \sum_{i \in \{0,1\}^n} \alpha_i |i\rangle \text{ with } \sum_{i \in \{0,1\}^n} |\alpha_i|^2 = 1.$$

Again,  $|\alpha_i|^2$  describes the probability of obtaining the basis state  $|i\rangle$  when measuring the qubits.

The state of qubits is altered using quantum operations, also known as quantum gates. Example of gates range from analogues of classical gates such as the X gate, which inverts the state of a qubit, to quantum-specific gates such as the Hadamard gate that can be used to put a qubit from a (classical) basis state into a (quantum) superposition. These gates are used to construct quantum circuits, the quantum computing analogue of classical circuits. Currently, quantum circuits are the predominant way to describe quantum computations.



**Figure 1: Five-qubit VQE quantum circuit. The circuit consists of chains of two-qubit CNOT gates, interleaved with single-qubit rotation gates.**

**EXAMPLE 1.** Fig. 1 shows a quantum circuit acting on five qubits. To this end, qubits are shown as horizontal lines and gates are denoted by boxes and symbols placed on these lines. Gates are applied in sequence from left to right. The circuit at hand consists of single-qubit rotation gates (indicated by purple boxes) and two-qubit CNOT gates (indicated in blue). It represents an instance of a popular near-term quantum algorithm (called Variational Quantum Eigensolver, or VQE [19]) applied to a small optimization problem.

### 2.2 Quantum Computing Architectures

Several technologies are currently being utilized or developed for quantum computing. Some of the most notable include superconducting qubits [20], trapped ions [21], neutral atoms [22], and photonic qubits [23]. Each of these technologies has distinct advantages and disadvantages. Superconducting qubits, for example, are very easy to control and manipulate, but they have to operate at extremely low temperatures. Trapped ions, on the other hand, can operate at much higher temperatures and have a long coherence time, but they are difficult to scale to a large number of qubits. It is crucial to note that most of the technologies are still in the research phase, and it is unclear which technology will prove to be most suitable for large-scale, general-purpose quantum computing. The remainder of this paper focuses on quantum computers based on superconducting qubits, as the quantum circuit mapping problem is most relevant for them.

**EXAMPLE 2.** Major industrial players such as IBM, Rigetti, and Google as well as academic institution such as Oxford Quantum Computing (OQC) currently provide access to quantum computers based on superconducting qubits. At the end of 2022, IBM, for example, provided devices ranging from 5 qubits all the way up to 433 qubits and they are planning for a system with more than 4000 qubits by 2025 in their roadmap ([ibm.com/quantum/roadmap](https://ibm.com/quantum/roadmap)).

Quantum computing architectures only offer a particular set of native gates, typically consisting of some single-qubit operations and a two-qubit entangling operation—effectively forming a universal gate set with which any quantum computation can be realized. This can be compared to a classical processor only supporting a particular Instruction Set Architecture (ISA). Just as a compiler is needed to convert high-level code to low-level instructions, high-level quantum circuits need to be decomposed to the native gate-set of the device they are supposed to be run on before being executable.

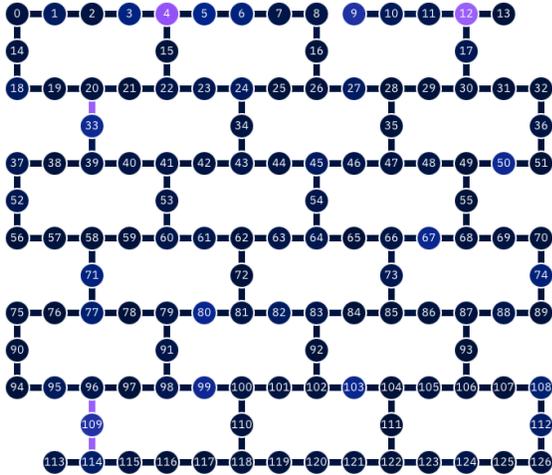


Figure 2: Coupling map of the 127-qubit *ibm\_washington* device

EXAMPLE 3. Again focusing on the quantum computers offered by IBM, they provide a native gate set that consists of three single-qubit gates ( $X$ ,  $\sqrt{X}$ , and  $R_z(\theta)$ ) as well as the two-qubit CNOT gate. Although seemingly small, this already constitutes a universal gate set for quantum computing, i.e., one out of which any conceivable quantum computation can be realized.

In addition, quantum computers, such as those based on superconducting qubits, only feature limited connectivity between their qubits, i.e., a two-qubit gate may not be applied to arbitrary qubits on the device. Usually, the ways in which qubits can interact with each other are described by a *coupling graph* (or *coupling map*). In general, this is a directed graph that describes which qubit pairs a certain two-qubit gate might be applied to. In most cases, the graph is actually undirected, meaning that the directionality of the gate between two qubits does not matter. To execute a circuit on one of these devices, the circuit has to respect the connectivity constraints imposed by the device’s coupling map.

EXAMPLE 4. Fig. 2 shows the coupling map of the 127-qubit *ibm\_washington* device. Like all recent IBM architectures, it is built from 12-qubit tiles that are stitched together in a brick-like pattern. Note how sparsely connected the overall architecture is, with each qubit having at most three neighbours.

### 3 QUANTUM CIRCUIT MAPPING AND MQT QMAP

As reviewed in the previous section, many quantum computing architectures limit the pairs of qubits that two-qubit operations can be applied to—commonly described by a device’s coupling map. To execute a generic quantum circuit (with arbitrary interactions between its qubits) on such an architecture, the circuit needs to be mapped. This involves *qubit allocation*, where the circuit’s (logical) qubits are assigned to the device’s (physical) qubits in an initial layout, and *routing*, where the original circuit is augmented with SWAP gates such that it adheres to the target device’s coupling map.

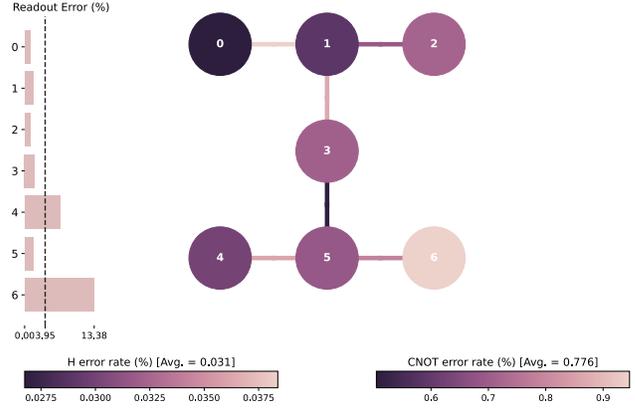


Figure 3: Seven-qubit *ibm\_oslo* architecture with device calibration data. Circles represent the individual qubits. Connections between qubits indicate that a two-qubit gate may be applied to them. The color of the qubits (edges) indicates the error rate of single-qubit (two-qubit) operations applied to that qubit (edge)—with darker implying lower error rates.

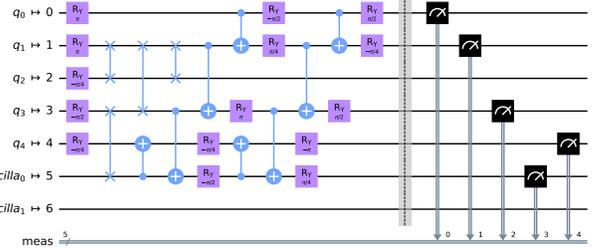
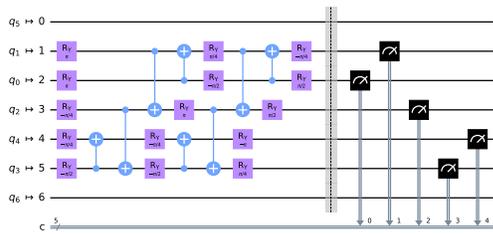


Figure 4: Naive mapping of the circuit from Fig. 1 to the *ibm\_oslo* architecture shown in Fig. 3. A total of four SWAP gates and an additional qubit are required to make the circuit executable.

EXAMPLE 5. Consider the circuit from Example 1 (shown in Fig. 1) and assume it shall be mapped to the seven-qubit *ibm\_oslo* architecture whose coupling map (together with some calibration information) is shown in Fig. 3. This circuit cannot be run directly on the architecture since it contains gates that act on qubits not connected on the device’s coupling map. Naively inserting SWAP gates that permute the logical-to-physical qubit mapping on the fly may yield a compiled circuit as shown in Fig. 4. Over the course of the mapping, four SWAP gates have been introduced to satisfy the connectivity constraints of the device’s architecture. Furthermore, the circuit uses an additional qubit for realizing the intended functionality. Since every additional gate (and qubit) increases the probability of errors, this is a very costly overhead for such a small circuit.



**Figure 5: Mapping of the circuit from Fig. 1 to the *ibm\_oslo* architecture shown in Fig. 3 using the exact mapper of QMAP. No SWAP gates or additional qubits are required to make the circuit executable.**

Keeping the number of additionally introduced gates as small as possible during quantum circuit mapping is key for ensuring the successful execution of the quantum circuit. Determining an optimal mapping for a quantum circuit is an NP-hard problem [24].

At the time of writing, the open-source quantum circuit mapping tool QMAP, which is part of the *Munich Quantum Toolkit* (MQT), offers two dedicated techniques for tackling that problem:

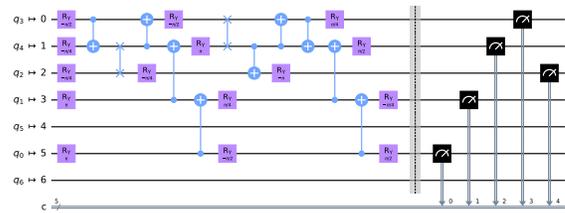
- (1) An exact mapping approach (based on [25], [26]) that guarantees (gate-optimal) solutions.
- (2) A heuristic mapping approach (based on [27]–[29]) that allows to determine efficient mapping solutions in a scalable fashion.

The exact mapper implemented in QMAP maps quantum circuits using the minimal number of SWAP gates. To this end, it encodes the mapping task as a MaxSAT problem and subsequently solves it using the SMT solver Z3 [30]. Due to the NP-hardness of the mapping task, this approach is only scalable up to roughly eight qubits and on the order of 1000 gates in most scenarios.

**EXAMPLE 6.** When used on the circuit shown in Fig. 1, the exact mapper of QMAP yields the circuit shown in Fig. 5. In this case, not a single SWAP gate and no additional qubits are required to map the circuit. Consequently, the expected performance of the resulting circuit is significantly improved. This performance comes at the cost of a mapping runtime in the order of a couple of seconds.

The heuristic mapper implemented in QMAP uses  $A^*$ -search [31] to efficiently traverse the immense search space of the mapping problem. It effectively trades optimality for runtime. This allows to reliably determine suitable mappings for circuits with up to hundreds of qubits and hundreds of thousands of gates.

**EXAMPLE 7.** When used on the circuit shown in Fig. 1, the heuristic mapper of QMAP yields the circuit shown in Fig. 6. While this solution is not optimal anymore it still only requires two SWAP gates and an additional CNOT gate as opposed to the four SWAP gates in the naive mapping. Furthermore, even for such a small example the heuristic mapper is orders of magnitudes faster than the exact mapper—requiring less than a millisecond.



**Figure 6: Mapping of the circuit from Fig. 1 to the *ibm\_oslo* architecture shown in Fig. 3 using the heuristic mapper of QMAP. The mapping uses two SWAP gates and one additional CNOT gate.**

## 4 USER’S PERSPECTIVE: WORKING WITH QMAP

Anyone may quickly and effectively map their quantum circuits to quantum computing systems with QMAP. The installation process and initial use of the MQT QMAP tool are briefly described in this section. By doing this, we offer a view into QMAP from the user’s perspective.

QMAP is mainly developed as a C++ library but, in order to make the tool as accessible and compatible as possible, comes with an easy-to-use Python interface. Installing the tool is as easy as

```
1 $ python3 -m venv .venv
2 $ source venv/bin/activate
3 (venv) $ pip install -U pip setuptools wheel
4 (venv) $ pip install mqt.qmap
```

In most practical cases (under 64-bit Linux, MacOS incl. Apple Silicon, and Windows), this requires no compilation and merely downloads and installs a platform-specific pre-built wheel.

The architecture to which a particular circuit should be mapped can be specified in a number of different ways. The simplest option is to just use one of IBM’s Qiskit’s [32] backends, but (arbitrary) architectures can also be created directly from a coupling map.

**EXAMPLE 8.** The *ibm\_oslo* device (shown in Fig. 3) considered previously can be instantiated via

```
1 from qiskit.providers.fake_provider import FakeOslo
2
3 arch = FakeOslo()
```

The same architecture could have also been provided without relying on Qiskit via

```
1 from mqt import qmap
2
3 arch = qmap.Architecture(7, {
4     (0, 1), (1, 0),
5     (1, 2), (2, 1),
6     (1, 3), (3, 1),
7     (2, 3), (3, 2),
8     (3, 5), (5, 3),
9     (4, 5), (5, 4),
10    (5, 6), (6, 5)}
11 )
```

With that, it is possible to define any desired quantum circuit and map it to the architecture. Once more, using Qiskit in Python is the easiest method. Alternative file formats, such as OpenQASM [33], also can be used to define the circuit.

EXAMPLE 9. The circuit shown in Fig. 1 originates from the MQT Bench library [34] which can be installed via

```
1 (venv) $ pip install mqt.bench
```

Then, the circuit can be obtained as a Qiskit QuantumCircuit via

```
1 from mqt import bench
2
3 qc = bench.get_benchmark(benchmark_name='vqe',
4                          circuit_size=5, level='indep')
```

Note that the particular parameters obtained for the rotation gates may vary with each invocation.

Once the architecture arch to map to and the quantum circuit qc to map have been defined, actually performing the mapping itself using QMAP only requires a single line of Python

```
1 qc_m, res = qmap.compile(qc, arch, method="exact")
```

for the exact mapper, or

```
1 qc_m, res = qmap.compile(qc, arch, method="heuristic")
```

for the heuristic mapper. Both return the mapped circuit as an IBM Qiskit QuantumCircuit object and a collection of results. The results, for example, allow one to conveniently print the required number of SWAP operations and the runtime of the mapping routine via

```
1 print("Additional SWAPs: %d" % res.output.swaps)
2 print("Runtime: %f" % res.time)
```

Naturally, the aforementioned walk-through only offered a brief overview of QMAP’s capabilities. The compile method has many optional arguments that can be used to configure the individual mappers. However, QMAP has comprehensive documentation that can be accessed at <https://mqtqmap.readthedocs.io> and offers a more extensive overview of all available options and techniques.

## 5 DEVELOPER’S PERSPECTIVE: IMPROVING QMAP

QMAP is a powerful tool, but it could be improved in many ways. Many aspects of QMAP were designed from a purely academic perspective with a rather narrow focus on very specific topics. This section begins with a quick overview of how to set up a development environment for QMAP, followed by a discussion of potential enhancement areas. More specifically, we show

- how the heuristic mapper could be made better by defining more complex initial layout algorithms as well as allowing for a wider variety of gates, and
- how the exact mapper’s SAT encoding and its construction could be adapted to improve the mapper’s overall performance.

This provides a brief look at QMAP from the developer’s perspective.

### 5.1 Setting up a Development Environment

QMAP is mainly developed as a C++ library that builds on the MQT Quantum Functionality Representation (QFR) library (<https://github.com/cda-tum/qfr>). In order to make the tool as accessible and compatible as possible, it comes with easy-to-use Python bindings. Consequently, developing features for the QMAP library mostly entails extending the core C++ library and, then, exposing the newly-added functionality to Python.

To get started, the project needs to be cloned from GitHub via

```
1 git clone --recursive https://github.com/cda-tum/qmap
```

Building the project requires a C++ compiler supporting C++17 and CMake with a minimum version of 3.19. In addition, the SMT solver Z3 [30] needs to be installed with a minimum version of 4.8.15 to build the exact mapper (see <https://github.com/Z3Prover/z3> for further details). Just as any other CMake-based project, building the C++ library is a two-step procedure

```
1 [.../qmap] $ cmake -S . -B build -DBUILD_QMAP_TESTS=ON
2 [.../qmap] $ cmake --build build
```

To locally build and install the Python package, first create a virtual environment as described in the previous section and, then, call

```
1 [.../qmap] (venv) $ pip install -e ".[dev]"
```

The library provides many convenience features, such as unit tests to test the library’s functionality, linting and code formatting tools to ensure proper quality of the code, and a comprehensive setup of CI workflows on GitHub that ensure the project builds and runs on all major operating systems. For further information, visit [mqtqmap.readthedocs.io/en/latest/DevelopmentGuide.html](https://mqtqmap.readthedocs.io/en/latest/DevelopmentGuide.html).

Having the development environment set up, QMAP can easily be extended and, by this, enhanced. In the following, examples for both, the heuristic mapper and the exact mapper are provided.

### 5.2 Improving the Heuristic Mapper

One direction how the heuristic mapper can be improved involves the initial mapping. As most quantum circuit mappers, the heuristic QMAP mapper splits the mapping tasks into (1) determining an initial layout for the placement of the logical qubits on the device’s physical qubits and (2) performing the actual routing based on the chosen initial layout. At the time of writing, QMAP supports three different kinds of initial layouts:

- (1) A trivial *identity* layout, where logical qubit  $q_i$  is assigned to physical qubit  $Q_i$  for all  $i$  from 0 to  $n - 1$ .
- (2) A *static* layout, which considers qubits that share a gate in the first layer of the circuit and maps those to any free connected qubit pair in the architecture while mapping the remaining qubits by order of index.
- (3) A *dynamic* layout, where qubits are greedily assigned on-demand when encountering them during the mapping.

As evaluations in the corresponding literature confirm [27]–[29], these initial layout techniques in combination with the  $A^*$  search already allow for rather efficient mapping results. At the same time, it is well known that the initial mapping has a significant impact on the routing performance that can be achieved; see, e.g., [35]. Consequently, more sophisticated initial mappings may lead to even better overall results—offering room for improvement. Many tools for mapping quantum circuits today (such as [36]) determine an initial mapping by focusing on matching a graph of the circuit’s gate interactions to the coupling map of the target architecture. This can be supplemented with device calibration information to account for differences in qubit and gate quality. Implementing such tactics could be a straightforward means of improving QMAP’s efficiency even further.

Another direction could involve the supported gate sets. For historical reasons, QMAP assumes that the only two-qubit gates in the circuit are CNOT gates. This was a reasonable assumption at a time when only IBM provided public access to their devices and Qiskit was the only major open-source quantum SDK. In the meantime, however, many further (and different) hardware providers and quantum computers have emerged—each with their own native gate set. As an example, Rigetti quantum computers expose the CZ gate rather than the CNOT gate. As a result, adapting QMAP to newer types of devices and/or to support more diverse gate sets is another straightforward means of improving QMAP.

### 5.3 Improving the Exact Mapper

Also the exact mapper offers room for improvements. As previously discussed, the exact mapper encodes the mapping task as a MaxSAT problem, allowing it to deliver results with the fewest possible SWAP gates. The quality of the results comes at a cost, however, because the search space of the mapping problem scales as  $N! * k$ , where  $N$  is the number of qubits in the architecture and  $k$  is the number of (two-qubit) gates in the circuit. As a result, the technique is currently limited to architectures with fewer than ten qubits and circuits with a few hundred to a thousand gates.

A number of efforts have been made to prune the search space in the quantum circuit mapping problem. It has been demonstrated in [26] that it is sufficient to consider just enough permutations in front of each layer in the mapping for any two qubits to be placed adjacently. This drastically reduces the number of permutations from  $N!$ . Nonetheless, the concepts described in [26] are implemented in a rather straightforward manner. Further improvements could be made in this area by streamlining the computation of the permutations that must be considered, for example, by developing a specialized Cayley graph datastructure and corresponding manipulation algorithms.

In addition, attempts have been made to limit the number of qubits that must be considered during the mapping. Without such restrictions, it would be impossible to map a 3-qubit circuit to a 100-qubit architecture, for instance. The most radical approach is given in [25], which considers all viable subarchitectures containing as many qubits as the circuit to be mapped rather than the entire architecture. While this reduces the number of qubits to the greatest extent possible (a circuit can never be mapped to an architecture with fewer qubits than the circuit itself), it is not without its own downsides. In the first place, the existing implementation is relatively simple, as it does not account for subarchitecture isomorphism, i.e., the possibility that many subarchitectures are structurally identical. Instead, it simply explores every possible combination using brute force. Furthermore, it has been demonstrated in [37] that the optimality of the resulting circuit cannot be guaranteed because subarchitectures with more qubits may permit a better mapping. Fortunately, [37] proposes methods for determining optimal or near-optimal sets of subarchitectures to be considered during mapping. These methods have not yet been incorporated directly into the exact mapper.

## 6 CONCLUSIONS

This summary paper described the MQT QMAP tool, which allows researchers and developers to effectively map quantum circuits to real quantum computing architectures. We briefly reviewed the quantum circuit mapping problem and demonstrated how QMAP can be used to solve this problem efficiently and in an automated fashion. Furthermore, we gave a rough overview of how to develop for QMAP and offered potential avenues for future improvement. In doing so, we hope to pique more interest in this intriguing and rapidly growing area of research. In that case, we are referring to the aforementioned references for a more in-depth discussion of the relevant methods and challenges.

### Acknowledgements

We thank everyone that contributed to the QMAP tool. Special thanks to Stefan Hillmich, Tom Peham, Sarah Schneider, and Alwin Zulehner for their specific contributions to the work presented here.

This work received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 101001318), was part of the Munich Quantum Valley, which is supported by the Bavarian state government with funds from the Hightech Agenda Bayern Plus, and has been supported by the BMWK on the basis of a decision by the German Bundestag through project QuaST.

### REFERENCES

- [1] Michael A. Nielsen and Isaac L. Chuang, *Quantum Computation and Quantum Information*. Cambridge University Press, 2010.
- [2] Peter W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM J. Comput.*, 1997.
- [3] Jacob Biamonte *et al.*, "Quantum machine learning," *Nature*, vol. 549, no. 7671, pp. 195–202, 2017. arXiv: [1611.09347](https://arxiv.org/abs/1611.09347).
- [4] Stuart Harwood *et al.*, "Formulating and solving routing problems on quantum computers," *IEEE Transactions on Quantum Engineering*, vol. 2, pp. 1–17, 2021. DOI: [10.1109/TQE.2021.3049230](https://doi.org/10.1109/TQE.2021.3049230).
- [5] Katherine L. Brown, William J. Munro, and Vivien M. Kendon, "Using quantum computers for quantum simulation," *Entropy*, vol. 12, no. 11, pp. 2268–2307, 2010. DOI: [10.3390/e12112268](https://doi.org/10.3390/e12112268).
- [6] Y. Cao, J. Romero, and A. Aspuru-Guzik, "Potential of quantum computing for drug discovery," *IBM J. Res. & Dev.*, vol. 62, no. 6, 6:1–6:20, 2018. DOI: [10.1147/JRD.2018.2888987](https://doi.org/10.1147/JRD.2018.2888987).
- [7] Benchen Huang, Marco Govoni, and Giulia Galli, "Simulating the Electronic Structure of Spin Defects on Quantum Computers," *PRX Quantum*, vol. 3, no. 1, p. 010339, 2022. DOI: [10.1103/PRXQuantum.3.010339](https://doi.org/10.1103/PRXQuantum.3.010339).
- [8] Vedran Dunjko and Hans J Briegel, "Machine learning & artificial intelligence in the quantum domain: A review of recent progress," *Rep. Prog. Phys.*, vol. 81, no. 7, p. 074001, 2018. DOI: [10.1088/1361-6633/aab406](https://doi.org/10.1088/1361-6633/aab406).

- [9] Wakaki Hattori and Shigeru Yamashita, “Quantum circuit optimization by changing the gate order for 2D nearest neighbor architectures,” in *Int’l Conf. of Reversible Computation*, vol. 11106, 2018, pp. 228–243. doi: [10.1007/978-3-319-99498-7\\_16](https://doi.org/10.1007/978-3-319-99498-7_16).
- [10] Zahra Sasanian and D. Michael Miller, “Reversible and quantum circuit optimization: A functional approach,” in *Int’l Conf. of Reversible Computation*, red. by David Hutchison *et al.*, vol. 7581, 2013, pp. 112–124. doi: [10.1007/978-3-642-36315-3\\_9](https://doi.org/10.1007/978-3-642-36315-3_9).
- [11] Toshinari Itoko, Rudy Raymond, Takashi Imamichi, and Atsushi Matsuo, “Optimization of quantum circuit mapping using gate transformation and commutation,” *Integration*, vol. 70, pp. 43–50, 2020. doi: [10.1016/j.vlsi.2019.10.004](https://doi.org/10.1016/j.vlsi.2019.10.004).
- [12] D. Maslov, G.W. Dueck, D.M. Miller, and C. Negrevergne, “Quantum circuit simplification and level compaction,” *IEEE Trans. on CAD of Integrated Circuits and Systems*, vol. 27, no. 3, pp. 436–444, 2008. doi: [10.1109/TCAD.2007.911334](https://doi.org/10.1109/TCAD.2007.911334).
- [13] Sarah Schneider, Lukas Burgholzer, and Robert Wille, “A SAT encoding for optimal Clifford circuit synthesis,” in *Asia and South Pacific Design Automation Conf.*, 2023.
- [14] Ed Younis and Costin Iancu, “Quantum circuit optimization and transpilation via parameterized circuit instantiation,” in *qce*, IEEE, 2022, pp. 465–475. doi: [10.1109/QCE53715.2022.00068](https://doi.org/10.1109/QCE53715.2022.00068).
- [15] G. Vidal and C. M. Dawson, “Universal quantum circuit for two-qubit transformations with three controlled-NOT gates,” *Phys. Rev. A*, vol. 69, no. 1, p. 010301, 2004. doi: [10.1103/PhysRevA.69.010301](https://doi.org/10.1103/PhysRevA.69.010301).
- [16] Dmitri Maslov, “On the advantages of using relative phase Toffolis with an application to multiple control Toffoli optimization,” *Phys. Rev. A*, vol. 93, no. 2, p. 022311, 2016. doi: [10.1103/PhysRevA.93.022311](https://doi.org/10.1103/PhysRevA.93.022311).
- [17] D. Michael Miller, Robert Wille, and Zahra Sasanian, “Elementary quantum gate realizations for multiple-control Toffoli gates,” in *Int’l Symp. on Multi-Valued Logic*, 2011. doi: [10.1109/ISMVL.2011.54](https://doi.org/10.1109/ISMVL.2011.54).
- [18] Marc G. Davis *et al.*, “Towards Optimal Topology Aware Quantum Circuit Synthesis,” in *qce*, IEEE, 2020, pp. 223–234. doi: [10.1109/QCE49297.2020.00036](https://doi.org/10.1109/QCE49297.2020.00036).
- [19] M. Cerezo *et al.*, *Variational quantum algorithms*, 2020. arXiv: [2012.09265](https://arxiv.org/abs/2012.09265).
- [20] P. Krantz *et al.*, “A quantum engineer’s guide to superconducting qubits,” *Applied Physics Reviews*, vol. 6, no. 2, p. 021318, 2019. doi: [10.1063/1.5089550](https://doi.org/10.1063/1.5089550).
- [21] H Haffner, C Roos, and R Blatt, “Quantum computing with trapped ions,” *Physics Reports*, vol. 469, no. 4, pp. 155–203, 2008. doi: [10.1016/j.physrep.2008.09.003](https://doi.org/10.1016/j.physrep.2008.09.003).
- [22] Loïc Henriët *et al.*, “Quantum computing with neutral atoms,” *Quantum*, vol. 4, p. 327, 2020. doi: [10.22331/q-2020-09-21-327](https://doi.org/10.22331/q-2020-09-21-327).
- [23] Jeremy L. O’Brien, Akira Furusawa, and Jelena Vučković, “Photonic quantum technologies,” *Nature Photon*, vol. 3, no. 12, pp. 687–695, 2009. doi: [10.1038/nphoton.2009.229](https://doi.org/10.1038/nphoton.2009.229).
- [24] A. Botea, A. Kishimoto, and Radu Marinescu, “On the complexity of quantum circuit compilation,” in *Int’l Symp Comb. Search*, 2018.
- [25] Robert Wille, Lukas Burgholzer, and Alwin Zulehner, “Mapping quantum circuits to IBM QX architectures using the minimal number of SWAP and H operations,” in *Design Automation Conf.*, 2019.
- [26] Lukas Burgholzer, Sarah Schneider, and Robert Wille, “Limiting the search space in optimal quantum circuit mapping,” in *Asia and South Pacific Design Automation Conf.*, 2022.
- [27] Alwin Zulehner, Alexandru Paler, and Robert Wille, “An efficient methodology for mapping quantum circuits to the IBM QX architectures,” *IEEE Trans. on CAD of Integrated Circuits and Systems*, 2019.
- [28] Stefan Hillmich, Alwin Zulehner, and Robert Wille, “Exploiting Quantum Teleportation in Quantum Circuit Mapping,” in *Design, Automation and Test in Europe*, ACM, 2021, pp. 792–797. doi: [10.1145/3394885.3431604](https://doi.org/10.1145/3394885.3431604).
- [29] Alwin Zulehner and Robert Wille, “Compiling SU(4) quantum circuits to IBM QX architectures,” in *Asia and South Pacific Design Automation Conf.*, 2019, pp. 185–190. doi: [10.1145/3287624.3287704](https://doi.org/10.1145/3287624.3287704).
- [30] Leonardo de Moura and Nikolaj Björner, “Z3: An efficient SMT solver,” in *Tools Algorithms Constr. Anal. Syst.*, Springer, 2008, pp. 337–340.
- [31] Peter Hart, Nils Nilsson, and Bertram Raphael, “A Formal Basis for the Heuristic Determination of Minimum Cost Paths,” *IEEE Trans. Syst. Sci. Cyber.*, vol. 4, no. 2, pp. 100–107, 1968. doi: [10.1109/TSSC.1968.300136](https://doi.org/10.1109/TSSC.1968.300136).
- [32] Matthew Treinish *et al.*, *Qiskit/qiskit*, Zenodo, 2022. doi: [10.5281/ZENODO.2573505](https://doi.org/10.5281/ZENODO.2573505).
- [33] Andrew W. Cross *et al.*, *OpenQASM 3: A broader and deeper quantum assembly language*, 2021. arXiv: [2104.14722](https://arxiv.org/abs/2104.14722).
- [34] Nils Quetschlich, Lukas Burgholzer, and Robert Wille, *MQT Bench: Benchmarking software and design automation tools for quantum computing*, 2022. arXiv: [2204.13719](https://arxiv.org/abs/2204.13719).
- [35] Alexandru Paler, “On the Influence of Initial Qubit Placement During NISQ Circuit Compilation,” in *Quantum Technol. Optim. Probl.*, vol. 11413, 2019, pp. 207–217. doi: [10.1007/978-3-030-14082-3\\_18](https://doi.org/10.1007/978-3-030-14082-3_18).
- [36] Seyon Sivarajah *et al.*, “T|ket): A retargetable compiler for NISQ devices,” *Quantum Sci. Technol.*, vol. 6, no. 1, p. 014003, 2021. doi: [10.1088/2058-9565/ab8e92](https://doi.org/10.1088/2058-9565/ab8e92).
- [37] Tom Peham, Lukas Burgholzer, and Robert Wille, *On Optimal Subarchitectures for Quantum Circuit Mapping*, 2022. arXiv: [2210.09321](https://arxiv.org/abs/2210.09321) [quant-ph].