

Alternating ZX Circuit Extraction for Hardware-Adaptive Compilation

Ludwig Schmid^{*1}, Korbinian Staudacher^{†1}, and Robert Wille^{*‡§}

¹These authors contributed equally to this work.

^{*}Chair for Design Automation, Technical University of Munich, Germany

[†]MNM-Team, Ludwig-Maximilians-University Munich, Germany,

[‡]Software Competence Center Hagenberg GmbH, Austria

[§]Munich Quantum Software Company GmbH, Germany

ludwig.s.schmid@tum.de; korbinian.staudacher@nm.ifi.lmu.de; robert.wille@tum.de

<https://www.cda.cit.tum.de/research/quantum/>

Abstract—We present a novel quantum circuit extraction scheme that tightly integrates graph-like ZX diagrams with hardware-adaptive routing. The method utilizes the degrees of freedom during the conversion from a ZX diagram to a quantum circuit (*extraction*). It alternates between generating multiple extraction options and evaluating them based on hardware constraints, allowing the routing algorithm to inform and guide the extraction process. This feedback loop extends existing graph-like ZX extraction and supports modular integration of different extraction algorithms, routing strategies, and target hardware, making it a versatile building block during compilation. To perform numerical evaluations, a reference instance of the scheme is implemented with SWAP-based routing for neutral atom hardware and evaluated using various benchmark collections on small- to mid-scale circuits. The reference code is available as open-source, allowing fast integration of other extraction and/or routing tools to stimulate further research and foster improvements of the proposed scheme.

I. INTRODUCTION

Executing quantum algorithms on real quantum hardware necessitates several preprocessing steps to translate general quantum operations into an executable, hardware-dependent form. As different hardware platforms exhibit different computational capabilities and operational constraints, this requires specialized optimization routines, generally referred to as *compilation* [1], [2]. Essential steps in this process include the generation and optimization of quantum circuits themselves, which we refer to as *circuit synthesis*, and the adaption of the synthesized circuit to a given hardware architecture, including qubit *routing*. Typically, these compilation steps are performed independently, with the output of one step serving as the input for the next. However, this sequential approach can yield suboptimal results, as the output of one step may not be ideal for the subsequent step [1].

For the circuit synthesis step, ZX calculus [3] has established itself as a powerful tool for designing [3]–[5] and optimizing [3], [6]–[8] quantum circuits. It involves converting a quantum circuit into a ZX diagram, simplifying it using diagrammatic rules and, then, converting it back into a quantum circuit through *circuit extraction* [9], [10]. This procedure can also be seen as transforming quantum circuits to measurement patterns, performing equivalence-preserving rewrites of the pattern, and extracting the pattern back into a quantum circuit [11], [12].

Existing circuit optimization algorithms using this pipeline usually operate on a hardware-agnostic level, e.g., by optimizing T-gate [6] or two-qubit gate count [7], [9]. Yet, the best possible circuit may heavily depend on the hardware and the employed routing strategy. In this work, we propose an approach for hardware-adaptive compilation based on the ZX calculus, which iteratively alternates between the circuit extraction and the hardware routing step to guide the extraction process. We exploit that the extraction step is not unique and gives rise to a plethora of possibilities, which we call *extraction paths*. These paths are evaluated based on the current hardware configuration, and the best path is selected to continue the extraction.

Advantages of this approach include:

- 1) Hardware-adaptive circuit extraction beyond simplistic metrics such as gate count.
- 2) Arbitrary hardware platforms can be targeted by employing different routing strategies.
- 3) Its simplicity allows for easy integration and adaptation of already existing extraction and routing software, including e.g., future fault-tolerant compilation tools.

The general scheme and its functionality are described, and multiple possible extensions to improve the performance of the base version are discussed. A reference implementation based on the default ZX extraction algorithm in combination with a SWAP-based routing algorithm for neutral atom hardware is provided. This reference implementation is evaluated using random circuits and different benchmarking sets, demonstrating fidelity improvements for the final output circuit. The results are discussed, focusing on potential future work.

The remainder of this paper is structured as follows: Sec. II provides a brief review of the ZX calculus and the extraction of graph-like ZX diagrams. Sec. III discusses the process of routing circuits to different hardware platforms. Then, Sec. IV illustrates the issue of suboptimal compilation due to the independent use of extraction and routing algorithms, followed by the proposed scheme in Sec. V. In Sec. VI, the scheme is evaluated, and the results are compared to the default extraction. Finally, we conclude the work in Sec. VII and discuss future work.




Figure 1. Definition of Z(X)-spiders and the identity (Hadamard) wire.




Figure 2. Two of the ZX-calculus rules, namely the fusion rule (f) merging connected spiders and the Hadamard rule (h) inverting colors.

II. ZX-CALCULUS PRELIMINARIES

The following section gives a brief overview of the ZX calculus and the extraction of quantum circuits from graph-like ZX diagrams. For a comprehensive introduction to the ZX-calculus, we refer to Ref. [3].

A. ZX-Calculus

The ZX-calculus is a diagrammatic language for reasoning about linear maps in quantum computing, where nodes (spiders) and edges (wires) form an undirected graph called a ZX diagram. There are two types of spiders: green *Z-spiders* and red *X-spiders*. Spiders can be parametrized with an angle $\alpha \in [0, 2\pi)$ and correspond to two-dimensional matrices in the Hilbert space with the definition given in Figure 1. In general, they can have any number of incoming and outgoing wires. For convenience, we distinguish between two types of wires: *Normal wires*, representing the identity, and *Hadamard wires*, representing the Hadamard matrix. Wires entering the diagram from the left (right) are called *input wires* (*output wires*), with the adjacent spiders defined as *inputs* I (*outputs* O).

ZX-diagrams constitute a graphical language for quantum circuits, meaning any quantum circuit can be represented as a ZX-diagram by replacing gates with equivalent diagrams. The power of the ZX-calculus lies in its rules for manipulating these diagrams. While not changing the underlying linear map, they can be used to simplify the diagram and, therefore, possibly the resulting quantum circuit.

Two rules used in the following are shown in Figure 2, namely the fusion rule (f) and the Hadamard rule (h). The fusion rule (f) allows merging spiders of the same color if they are connected by at least one normal wire, and (h) allows changing the colors of spiders by flipping normal and Hadamard wires. All rules apply in both directions and remain valid with interchanged colors. These two rules can be extended to form a complete graphical rule set [13].

Example 1. Figure 3 shows how a two-qubit Grover search circuit is translated to a ZX-diagram, replacing also the two-qubit gates with their equivalent ZX-diagram represen-




Figure 3. Example of a two-qubit Grover search circuit, which is translated first to a ZX diagram and then converted to a graph-like ZX-diagram using the rules from Figure 2.




Figure 4. Basic circuit extraction from a graph-like ZX diagram.

tation. The diagram is then further simplified using the fusion and the Hadamard rule. *

B. Graph-like ZX-diagrams

In this work, we focus on the special class of *graph-like* ZX-diagrams as introduced in Ref. [11], which allows us to represent any quantum computation as a graph of parametrized green Z-spiders and Hadamard wires. One can transform any ZX-diagram into an equivalent graph-like ZX-diagram by repeatedly applying standard ZX-rules [11].

Example 2. The last step in Figure 3 shows the graph-like ZX-diagram of the previous two-qubit Grover search circuit. *

Graph-like diagrams can be directly interpreted as patterns in the measurement-based quantum model where each spider corresponds to a qubit in $|+\rangle$ state and each Hadamard wire to a CZ gate. As such, they provide additional transformations such as *local complementation*, *pivoting*, and *phase gadget elimination*, where the latter was introduced in Ref. [6]. Using these rules to simplify graph-like diagrams has become a default procedure and is offered by tools such as PYZX [14].

C. Circuit Extraction

While the diagrammatic language is a powerful tool for reasoning and optimizing quantum operations, as a final step, to obtain a quantum circuit with the same number of qubits as the original, one can *extract* a quantum circuit from the ZX diagram. This extraction is feasible in polynomial time if the underlying graph has some kind of *flow* [9], [15]. Here, we provide a brief overview of the extraction algorithm for graph-like ZX diagrams with gflow as described in Ref. [9].

The algorithm extracts a quantum circuit by identifying suitable parts of the ZX diagram and converting them into their equivalent quantum gates. These parts are then removed from the diagram, extracting one gate at a time until only the inputs and outputs remain. During the process, a set of green




Figure 5. Full ZX circuit extraction of a graph-like diagram using a CX gate to rearrange the graph connections.

Z-spiders called the *frontier* separates the extracted portion of the diagram from the unextracted part.

First, there are three basic extraction rules:

- **Phase:** Phases of frontier spiders are directly extracted as R_z gates.
- **CZ:** Hadamard wires between frontier spiders are extracted as CZ gates.
- **Hadamard:** Hadamard wires where a frontier spider is only connected to a non-frontier spider are extracted as Hadamard gates.

Example 3. Figure 4 illustrates the three above rules, extracting four gates from the diagram. The double vertical lines visualize the separation between the diagram and the circuit. *

Second, if every spider in the frontier has at least two non-frontier neighbors, one can add wires of a frontier spider to another by placing a CX gate in the extracted circuit with the following effect:

- 1) **CX:** Extracting a CX gate with control c and target qubit t copies all Hadamard wires of the target frontier spider to the control qubit frontier spider. Double wires cancel each other.

Example 4. In the upper part (i) of Figure 5, a CX gate between control qubit q_0 and target q_1 is extracted. The two Hadamard wires of the frontier spider q_1 are copied to q_0 and cancel the existing Hadamard wires. After this CX gate, the rest of the diagram can be extracted, as shown on the right. *

Previous work focused on finding optimized CX sequences and graph manipulations to reduce hardware requirements, such as two-qubit count [7], [9]. In addition, there exist special-purpose extraction algorithms, such as for multi-controlled phase gates [10]. In the following, we discuss how such extracted circuits can then be executed on hardware.

III. HARDWARE MAPPING & ROUTING

While the ZX calculus is a powerful tool to construct optimized quantum circuits, referred to as circuit *synthesis*, to execute these circuits on real hardware, further compilation steps are required [1], [2]. First *mapping*, which corresponds to a bijective assignment between the circuit qubits and the available hardware qubits. Second *routing*, where the limited hardware qubit connectivity is circumvented by inserting additional operations. Depending on the hardware platform

and the exact hardware specification, this step can be realized differently. This encompasses the insertion of SWAP gates [16]–[19], shuttling operations on QCCD trapped-ion architectures [20]–[23], or atom reconfigurations on the neutral atom platform [24]–[28]. A comprehensive discussion of the different methods and the corresponding software packages is out of the scope of this work, but the following are some general aspects of hardware compilation:

First, typically, the mapping and routing steps are executed after circuit synthesis. This is based on the promise that optimizing each step independently will yield a good overall result while facilitating the development of specialized algorithms for each step.

Second, optimization methods typically focus on some metric or proxy, such as (two-qubit) gate count. Yet, in all metrics, the overall cost function is the infidelity of the output circuit, which we aim to minimize in the mapping and routing process.

Third, the diversity of existing hardware, even within the same platform, requires highly specialized compilation methods to make optimal use of the available hardware. As a result, often only a small subset of the available compilation methods can be employed for a given hardware setup.

Within this work, we model the hardware compilation process as a black box, which takes a quantum circuit and a hardware architecture as inputs and returns a mapped quantum circuit together with some cost metrics indicating the overhead introduced by the mapping and routing process.

IV. CONSIDERED PROBLEM

In the following, we illustrate how the disjoint use of ZX-based circuit synthesis and hardware routing can lead to suboptimal results. For that, let us revisit the circuit extraction from Example 4 shown in the upper part (i) of Figure 5.

The circuit was extracted using a single additional CX gate, resulting in a total of three entangling gates (2 CZ and 1 CX) requiring interaction between qubits $q_0 \leftrightarrow q_2$ and $q_0 \leftrightarrow q_1$.

But, as discussed in Sec. II-C, this extraction step is not unique, and there is a certain degree of freedom in the extraction process [11]. Alternatively, one could have inserted a second CX gate between q_2 and q_3 to further reduce the number of Hadamard wires in the graph. This step and its result are shown as (ii) below, where the resulting circuit has four two-qubit gates (2 CZ and 2 CX) and requires connectivity between qubits $q_0 \leftrightarrow q_1$ and $q_1 \leftrightarrow q_2$.

For circuit extraction, a typical metric is the number of entangling gates, as they are a major source of error on available hardware [29]. Thus, the first circuit (i) would be preferable as it requires three entangling gates compared to four in the second circuit (ii).

However, this changes if one simultaneously considers the necessary routing. As an example, we assume SWAP-based routing and a linear hardware configuration as given on the very right of Figure 5. While the second circuit can be executed directly on the hardware, as all required qubit connections ($q_0 \leftrightarrow q_1$ and $q_1 \leftrightarrow q_2$) are directly available, the first circuit contains an unavailable two-qubit gate connection ($q_0 \leftrightarrow q_2$). The necessary SWAP insertion between q_0 and q_1 would result in an additional 3 CX gates, resulting in $3+3 = 6$ two-qubit gates for the first circuit (i). This makes the second circuit (ii) the better choice with its four two-qubit gates, as it inherently matches the hardware connectivity.

This simple example illustrates that a circuit that might be preferable based on metrics, such as the number of two-qubit gates, eventually can result in more costly routing and, therefore, in suboptimal final output. Although the example only considered SWAP gate insertions, the idea also applies to other routing strategies, such as shuttling or atom reconfigurations.

To circumvent this suboptimal interplay, we propose an alternating scheme where the extraction and routing algorithms communicate iteratively to avoid compilation steps that might be suboptimal for the opposite party.

V. ALTERNATING CIRCUIT EXTRACTION

In this section, we present a solution to address the issue of suboptimal compilation results arising from the disjoint use of ZX circuit extraction and hardware routing. First, we give a high-level overview of the proposed solution, followed by a detailed description of the implementation and further refinements. Third, we briefly discuss a concrete instance to realize this general scheme in practice, which will be used in the evaluations of Sec. VI.

A. Overview

We propose a circuit extraction scheme alternating between the extraction of graph-like ZX diagrams and the hardware routing, which satisfies the given hardware constraints. The abstract idea is illustrated in Figure 6, revisiting the circuits from the previous examples. The extraction procedure consists of the following steps:

- (a) **Circuit Extraction:** Within the ZX extraction, different possible extraction paths are generated with Phase, CZ, Hadamard and CX extraction each applied once.
- (b) **Path Information:** Descriptions of the paths (their full circuit or certain metrics) are sent to the routing process.
- (c) **Hardware Evaluation:** The routing algorithm evaluates the paths based on the current hardware configuration and selects the best path based on a chosen metric.
- (d) **Feedback:** The selection is sent back to the extraction algorithm, which continues the extraction process along the selected path.

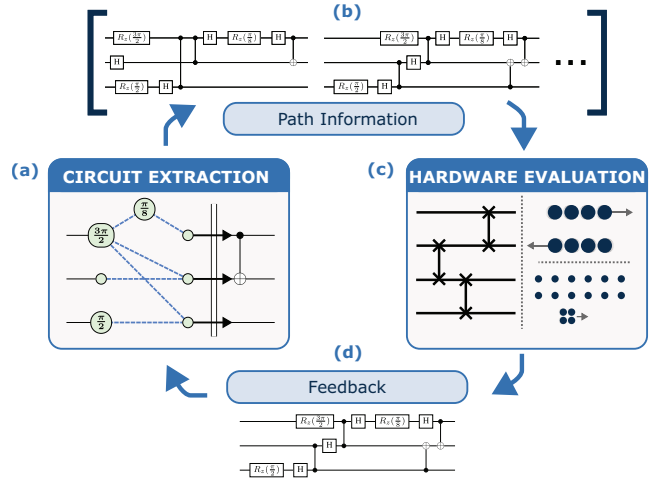


Figure 6. Overview of the cross-compilation strategy and the communication between circuit extraction and routing.

Through this iterative back-and-forth communication, the entire graph is converted into a fully mapped quantum circuit. It should be noted that this scheme is quite general and not limited to a particular extraction algorithm, evaluation metrics, routing strategy, or hardware platform. In particular, any of the extraction algorithms from Sec. II-C and any hardware compilation method from Sec. III could be employed. Furthermore, this could also be generalized to fault-tolerant compilation methods in a similar way.

While the feedback allows the extraction algorithm to make better-informed decisions, it also imposes an additional computational overhead for the iterative path evaluations. Let k denote the number of spiders in the diagram and m the number of paths at each extraction step. Since at each extraction step we remove at least one spider, we get an upper bound of $\mathcal{O}(km)$ evaluations of the routing cost. Assuming k to be constant for a given diagram, this results in a typical search-space exploration trade-off. By considering more extraction options, one will likely improve the output quality of the scheme while simultaneously increasing the necessary computational resources. Finding efficient heuristics for search space exploration and pruning constitutes, therefore, an essential task for future research to scale the scheme to larger systems. One should note that, given the resources, the costly evaluation of different extraction paths can be done in parallel.

B. Improving the Base Scheme

Based on the above scheme, we introduce several possible modifications that can further improve the performance of the base version. The different modifications are then also evaluated in Sec. VI.

Edge Bias Weight β : While all extraction paths are valid options, they are not necessarily equivalent regarding the “amount of graph” that is extracted. Simultaneously, shorter circuits are typically easier to route. As a result, the hardware routing might prefer shorter circuits over paths, which would result in a faster extraction. To account for this, we introduce

a cost bias with weight β to the evaluation metric, namely

$$C_{\text{bias}} = \beta \cdot \Delta W, \quad (1)$$

where ΔW is the reduction of Hadamard wires in the graph for a certain extraction path. Increasing β will favor extracting larger parts of the graph while potentially overshadowing the actual routing cost.

Sliding Window s : Many routing algorithms make heavy use of contextual information, including a certain lookahead to find globally favorable routing operations [16], [17], [27]. Considering every extraction circuit by itself neglects the surrounding circuit, resulting in locally optimized but globally suboptimal routing. This drawback affects both the evaluations of the different paths and the routing of the final output.

To circumvent this, we introduce a sliding window approach. Given the size of the window s , the routing algorithm will consider not only the current path but also the s previously extracted gates. For s sufficiently large, this results in a complete rerouting of the full extracted circuit. While giving more context to the routing process, it introduces additional overhead for each of the $\mathcal{O}(km)$ path evaluations.

Extraction Depth l : To provide even more context to the routing algorithm, one can also increase the extraction depth l which determines how often we run a cycle of Phase, Hadamard, CZ and CX extraction before hardware evaluation. This will provide the routing algorithm with all possible paths for l consecutive extractions. While providing more information to the routing algorithm, this exponentially increases the number of paths to be evaluated to $\mathcal{O}((km)^l)$, making it practically feasible only for small values of l .

VI. EVALUATIONS

While the scheme itself is general, we need to implement a concrete instance for numerical evaluations. For each of the parts in Figure 6, we use the following tools:

(a + b) For the ZX extraction phase, we modify the “default” extraction algorithm as described in Ref. [9] and send the complete circuits to the routing algorithm. For the representation of the graph-like ZX diagrams, we use `PyZX` [14].

(c + d) Hardware routing is realized using the neutral atom hybrid routing [27] from the Munich Quantum Toolkit (MQT [30]). While developed for neutral atom hardware, we only focus on the SWAP insertion part using the SABRE method [17], which could also be used for superconducting hardware. The best path is selected based on the *approximate success probability* (ASP), as defined in Ref. [31]

$$\text{ASP} = \exp\left(-\frac{t_{\text{idle}}}{T_{\text{eff}}}\right) \prod_{i=1}^N \mathcal{F}(O_i), \quad (2)$$

where $\mathcal{F}(O_i)$ represents the fidelity of the i -th operation, t_{idle} is the idle time, and $T_{\text{eff}} = \frac{T_1 T_2}{T_1 + T_2}$ denotes the effective coherence time. The ASP functions as a proxy to the actual fidelity expected on hardware by considering hardware parameters such as coherence times and gate fidelities.

The complete code of this reference implementation is available as open-source on Zenodo [32].




Figure 7. Change of fidelity compared to default extraction for different edge weight biases β on random Clifford+T circuits.




Figure 8. Change of fidelity compared to default extraction for different edge window sizes s and extraction depths l on random Clifford+T circuits.

A. Evaluation Setup

To study the performance and behavior of the proposed scheme, we perform numerical simulations using the setup described above. First, the influence of the different parameters, such as the extensions of Sec. V-B, are evaluated. Based on this, we apply the scheme to different benchmarking sets. Throughout this section, all circuits are converted to ZX diagrams and optimized using the `PyZX full_reduce` optimization. As a baseline, we employ the default `PyZX` extraction algorithm `extract_circuit` to generate the corresponding quantum circuit. We consider the ASP from Equation (2) of the final routed circuit as the target metric and calculate the relative difference compared to the baseline, referred to as Δ Fidelity. The hardware configuration is a 6×6 nearest neighbor grid with parameters taken from [31]. More details, including all evaluation scripts, data, and figures, are available on Zenodo [32].

B. Random Circuits

We consider random Clifford+T circuits generated using `PyZX randomT` function with $n = 6$ qubits and a circuit depth of $d = 300$. The T (CX) gate probability is set to $p_T = 0.4$ ($p_{CX} = 0.3$).

First, Figure 7 shows the result of varying edge weight bias β with the sliding window size set to the full circuit depth. For the considered circuits, the proposed scheme always improves the ASP of the output circuit. The data shows how the fidelity improvement increases for nonzero edge biases β , demonstrating the positive effect of including the number of extracted Hadamard wires into the cost function. For even larger biases, the performance drops slightly again, which may be due to the fact that the bias overshadows the actual routing cost, as discussed in Sec. V-B.

Second, Figure 8 shows the results of the other two introduced parameters, namely the varying window size s and




Figure 9. Change of fidelity compared to the default extraction for different benchmark sets. The base scheme (blue) refers to no edge bias $\beta = 0$, while the biased scheme (green) was chosen over a range of edge biases.

extraction depth l . As expected, it shows a performance increase for larger window sizes, supporting the hypothesis that providing more context to the routing algorithm is beneficial. It saturates at about $s \approx 2/3d$ of the total circuit depth. Similar output quality can be achieved by increasing the extraction depth l , which also facilitates the discovery of good extractions by providing more and deeper extraction information. Here, an edge bias of $\beta = 0.002$ was chosen.

C. Benchmarking Circuits

We evaluate circuits with qubit numbers ranging from 2 to 14 and circuit depths from 1 to 1337. The circuits are taken from common benchmarking libraries, including MQT Bench [33], QASM Bench [34], and the Feynman benchmarking set [35]. Figure 9 shows the relative difference of the ASP metric for different benchmarking sets, distinguished by different markers. The colors distinguish between the base scheme with $\beta = 0$ and a biased version, where the edge weight was optimized over different values $\beta \in \{5 \times 10^{-5}, 5 \times 10^{-4}, 1 \times 10^{-3}, 3 \times 10^{-3}, 5 \times 10^{-3}, 0.01\}$. The extraction depth was set to $l = 1$, and the window size was set to $s = \infty$, completely rerouting the whole circuit.

While for shallow circuits (typically preparation of certain states) the performance is similar to the default extraction, for deeper circuits, there can be large fidelity improvements of up to 250% for the biased version. At the same time, there also exist a few circuits where the default PyZX is slightly superior, even to the biased version. These circuits seem to have a specific structure that does not allow the routing evaluation to choose an overall beneficial extraction path. Future research should investigate this further and how to potentially circumvent this issue.

D. Discussion

The proposed setup performs well for the considered small and mid-sized circuits. However, as the employed routing algorithm relies heavily on the context of the circuit, it requires

repetitively rerouting large parts of the extracted circuit. In combination with the edge weight bias, the method requires specific hyperparameter optimization to outperform the default PyZX extraction. Nevertheless, the proposed scheme formulates a promising approach to realize hardware-adaptive ZX extraction. First, it extends the default ZX extraction for hardware-adaptive compilation beyond simple gate count metrics, where the original procedure is recovered if one considers only a single extraction path. Second, the structure is modular, where both extraction and routing can be replaced by other algorithms. This allows the scheme to be applied to arbitrary hardware platforms and their unique specialties while using existing compilation software. Third, its simplicity allows for fast integration of upcoming routing methods and adapting existing tools.

VII. CONCLUSION & OUTLOOK

We introduced an alternating scheme for quantum circuit extraction that combines ZX diagram extraction with hardware-adaptive routing through continuous feedback. By evaluating multiple extraction paths against hardware constraints and routing costs, the method enables better-informed extraction decisions and extends existing ZX extraction techniques. The modular structure allows for integration with various extraction algorithms and routing tools, making it applicable across different hardware architectures. It works on top of existing routing and compilation software, providing further improvements with little to no modifications.

This initial work opens several avenues for future research. First, one can consider additional extraction steps, for instance, the extraction of multi-controlled phase gates [10]. Second, the extension to alternative routing software, e.g., for trapped ion systems or even fault-tolerant routing. Finally, by investigating improved heuristics and search space pruning techniques, one should be able to reduce the number of paths to be evaluated while maintaining similar output quality.

ACKNOWLEDGMENTS

This work is partially supported by the German Federal Ministry of Education and Research (BMBF) under the funding program Quantum Technologies - From Basic Research to Market under contract number 13N16070. The authors acknowledge funding from the Munich Quantum Valley initiative (K5), which is supported by the Bavarian state government with funds from the Hightech Agenda Bayern Plus.

L.S. and R.W. acknowledge funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (Grant Agreement No. 101001318 and No. 101114305 -"MILLENION-SGA1"). Furthermore, this work was supported by the BMFTR under grant number 13N17298 (SYNQ) and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under grant numbers 563402549 and 563436708.

REFERENCES

- [1] F. T. Chong *et al.*, "Programming languages and compiler design for realistic quantum hardware," *Nature*, vol. 549, no. 7671, pp. 180–187, 2017.
- [2] J. Cong, "Lightning Talk: Scaling Up Quantum Compilation – Challenges and Opportunities," in *Design Automation Conf.*, 2023, pp. 1–2.
- [3] A. Kissinger and J. van de Wetering, *Picturing Quantum Software: An Introduction to the ZX-calculus and Quantum Compilation*. Preprint, 2024.
- [4] A. Kissinger, *Phase-free ZX diagrams are CSS codes (...or how to graphically grok the surface code)*, 2022. arXiv: 2204.14038.
- [5] H. Bombin *et al.*, "Unifying flavors of fault tolerance with the ZX calculus," *Quantum*, vol. 8, p. 1379, 2024.
- [6] A. Kissinger and J. van de Wetering, "Reducing the number of non-Clifford gates in quantum circuits," *Physical Review A*, vol. 102, no. 2, p. 022406, 2020.
- [7] K. Staudacher *et al.*, "Reducing 2-QuBit Gate Count for ZX-Calculus based Quantum Circuit Optimization," in *Int'l Conf. on Quantum Physics and Logic*, vol. 394, 2023, pp. 29–45.
- [8] V. Vandaele, *Qubit-count optimization using ZX-calculus*, 2024. arXiv: 2407.10171.
- [9] M. Backens *et al.*, "There and back again: A circuit extraction tale," *Quantum*, vol. 5, p. 421, 2021.
- [10] K. Staudacher *et al.*, "Multi-controlled phase gate synthesis with zx-calculus applied to neutral atom hardware," *Electronic Proceedings in Theoretical Computer Science*, vol. 406, pp. 96–116, 2024.
- [11] R. Duncan *et al.*, "Graph-theoretic Simplification of Quantum Circuits with the ZX-calculus," *Quantum*, vol. 4, p. 279, 2020.
- [12] P. Walther *et al.*, "Experimental one-way quantum computing," *Nature*, vol. 434, no. 7030, pp. 169–176, 2005.
- [13] R. Vilmart, "A Near-Optimal Axiomatisation of ZX-Calculus for Pure Qubit Quantum Mechanics," 2018.
- [14] A. Kissinger and J. van de Wetering, "PyZX: Large Scale Automated Diagrammatic Reasoning," in *Int'l Conf. on Quantum Physics and Logic*, vol. 318, 2020, pp. 229–241.
- [15] W. Simmons, "Relating measurement patterns to circuits via pauli flow," *Electronic Proceedings in Theoretical Computer Science*, vol. 343, pp. 50–101, 2021.
- [16] A. Zulehner *et al.*, "An Efficient Methodology for Mapping Quantum Circuits to the IBM QX Architectures," *IEEE Trans. on CAD of Integrated Circuits and Systems*, vol. 38, no. 7, pp. 1226–1236, 2019.
- [17] A. Cowtan *et al.*, "On the Qubit Routing Problem," in *Conference on the Theory of Quantum Computation, Communication and Cryptography*, vol. 135, 2019, 5:1–5:32. eprint: 1902.08091.
- [18] B. Tan and J. Cong, "Optimal Layout Synthesis for Quantum Computing," in *Int'l Conf. on CAD*, 2020, pp. 1–9.
- [19] G. Li *et al.*, "Tackling the Qubit Mapping Problem for NISQ-Era Quantum Devices," in *Int'l Conf. on Architectural Support for Programming Languages and Operating Systems*, 2019, pp. 1001–1014.
- [20] D. Schoenberger *et al.*, "Shuttling for Scalable Trapped-Ion Quantum Computers," *IEEE Trans. on CAD of Integrated Circuits and Systems*, pp. 1–1, 2024.
- [21] P. Murali *et al.*, "Architecting Noisy Intermediate-Scale Trapped Ion Quantum Computers," in *IEEE International Symposium on Computer Architecture*, 2020, pp. 529–542.
- [22] F. Kreppel *et al.*, "Quantum Circuit Compiler for a Shuttling-Based Trapped-Ion Quantum Computer," *Quantum*, vol. 7, p. 1176, 2023.
- [23] A. A. Saki *et al.*, "Muzzle the Shuttle: Efficient Compilation for Multi-Trap Trapped-Ion Quantum Computers," in *Design, Automation and Test in Europe*, 2022, pp. 322–327. arXiv: 2111.07961.
- [24] D. B. Tan *et al.*, "Compiling quantum circuits for dynamically field-programmable neutral atoms array processors," *Quantum*, vol. 8, p. 1281, 2024.
- [25] W.-H. Lin *et al.*, *Reuse-Aware Compilation for Zoned Quantum Architectures Based on Neutral Atoms*, 2024. arXiv: 2411.11784.
- [26] Y. Stade *et al.*, "An abstract model and efficient routing for logical entangling gates on zoned neutral atom architectures," in *Int'l Conf. on Quantum Computing and Engineering*, IEEE, 2024, pp. 784–795.
- [27] L. Schmid *et al.*, "Hybrid Circuit Mapping: Leveraging the Full Spectrum of Computational Capabilities of Neutral Atom Quantum Computers," in *Design Automation Conf.*, 2024, pp. 1–6.
- [28] J. Ludmir and T. Patel, "Parallax: A Compiler for Neutral Atom Quantum Computers under Hardware Constraints," in *Int'l Conf. for High Performance Computing*, 2024, pp. 1–17.
- [29] J. Preskill, "Quantum Computing in the NISQ era and beyond," *Quantum*, vol. 2, p. 79, 2018.
- [30] R. Wille *et al.*, "The MQT Handbook: A Summary of Design Automation Tools and Software for Quantum Computing," 2024. arXiv: 2405.17543.
- [31] L. Schmid *et al.*, "Computational capabilities and compiler development for neutral atom quantum processors—connecting tool developers and hardware experts," *Quantum Science and Technology*, vol. 9, no. 3, p. 033001, 2024.
- [32] L. Schmid *et al.*, *Alternating ZX Circuit Extraction for Hardware-Adaptive Compilation*, Zenodo, 2026. DOI: 10.5281/zenodo.18199412.
- [33] N. Quetschlich *et al.*, "MQT Bench: Benchmarking Software and Design Automation Tools for Quantum Computing," *Quantum*, vol. 7, p. 1062, 2023.
- [34] A. Li *et al.*, *QASMBench: A Low-level QASM Benchmark Suite for NISQ Evaluation and Simulation*, 2022. eprint: 2005.13018.
- [35] M. Amy, "Towards Large-scale Functional Verification of Universal Quantum Circuits," *Electronic Proceedings in Theoretical Computer Science*, vol. 287, pp. 1–21, 2019. eprint: 1805.06908.