Towards Supporting QIR: Thoughts on Adopting the Quantum Intermediate Representation

Yannick Stade*, Lukas Burgholzer*, and Robert Wille*†
*Chair for Design Automation, Technical University of Munich, Munich, Germany
†Software Competence Center Hagenberg GmbH, Hagenberg, Austria
{yannick.stade, lukas.burgholzer, robert.wille}@tum.de
www.cda.cit.tum.de/research/quantum

***Abstract*—New records in the number of qubits and the fidelity of quantum computers continue to be set. Additionally, the quantum computing community is eager to leverage this immense computational power. However, to execute an application on hardware, it has to be translated into a sequence of hardware-specific instructions. To this end, intermediate representations play a crucial role in the software stack for a quantum computer to facilitate efficient optimizations. One of those intermediate representations is the *Quantum Intermediate Representation* (QIR), proposed by MICROSOFT. In this article, we provide food for thought on how QIR can be adopted in different software tools. We discuss the advantages and disadvantages of various approaches and outline related challenges. Finally, we conclude with an outlook on future directions using QIR.**

***Index Terms*—quantum computing, intermediate representation, compiler, runtime, software stack**

I. INTRODUCTION

Since the famous results of Shor [1], the promise holds that quantum computers can offer an exponential speedup over classical computers for specific problems. Many applications [2]–[6] have since been developed to exploit this potential computational power [7], [8]. However, physical implementations of quantum computers still need to catch up and are far from being useful for the aforementioned applications, with only a few realizations of toy examples on real hardware devices reported, e. g., [9], [10]. Nevertheless, the development of real quantum computers has seen rapid progress over the past years. Breakthroughs in different technologies, such as superconducting qubits by companies like GOOGLE and IBM, trapped ions [11], and neutral atoms [10], [12], have demonstrated significant advancements in the number of qubits and their reliability.

However, hardware is only one part of the story: Successfully realizing applications on quantum computers requires a sophisticated stack of software tools to transform a classical problem description into a quantum solution, to optimize the resulting (hybrid) quantum-classical program, and, finally, to execute it [13], [14]. Since real quantum computers are not capable of executing arbitrary instructions directly, a compiler with various transformation and optimization passes is needed [15]–[19]. These passes are essential to minimize the overhead introduced by the compilation process and to maintain a high fidelity of the resulting quantum program. Many of the problems involved in this process are computationally hard to solve, making the development of such a compiler or a general compiler infrastructure challenging.

Intermediate representations are a key component of any compiler, enabling the seamless combination of various passes and transformations. Huge parts of the quantum computing ecosystem have come to rely on the OpenQASM format, proposed initially as version 2 [20] and refined in version 3 [21], by IBM. The format started as a low-level quantum assembly language and has since evolved into a more general-purpose quantum programming language adding classical elements.

In contrast, the *Quantum Intermediate Representation* (QIR), initially proposed by MICROSOFT [22], takes a different approach. It aims to reuse as much as possible of the existing classical compiler infrastructure, namely the LLVM compiler framework [23], and augment it with quantum instructions. Despite its promises, the adoption of QIR in the quantum computing ecosystem remains low, especially outside the realm of industrial players like MICROSOFT, NVIDIA, or XANADU that have the resources and the expertise to develop and maintain such a compiler infrastructure.

In this overview article, we aim to shine a light on the challenges and opportunities that come with adopting QIR in the quantum computing software ecosystem. This should help readers to develop a better understanding of the various options to support QIR in their software tools.

Throughout the article, we assume the reader to be familiar with the basics of quantum computing and quantum circuits. We refer the reader to the literature [24] for an in-depth introduction.

II. AN ABRIDGED HISTORY OF QUANTUM INTERMEDIATE REPRESENTATIONS

First, this section briefly reviews the history of quantum intermediate representations and how quantum computations are typically represented. While certainly not exhaustive, these recollections provide a starting point for understanding the different approaches to quantum intermediate representations.

A. OpenQASM 2: Quantum Computation $\approx$ Quantum Circuit

The *Open Quantum Assembly* (OpenQASM) is the language initially proposed in version 2 by IBM in 2017 [20] to describe quantum experiments to be executed on their publicly available quantum computers. It is a low-level language that describes quantum computations in a straightforward manner, i. e., as an enumeration of quantum instructions. Additionally, it supports measurements, resets, (limited) feedback, and gate subroutines. Since its inception, OpenQASM has been widely adopted throughout the quantum computing community and is heavily used as an exchange format for quantum circuits.

Example 1. *The "Hello World" of quantum computing, the creation of a Bell state, is expressed in OpenQASM 2.0 on the top left in Fig. 1. After loading the quantum (standard) library, the code declares a quantum register with two qubits and a classical register with two bits. Then, an Hadamard-gate is applied to the first qubit, followed by a CNOT-gate controlled by the first qubit and targeting the second qubit. Finally, both qubits are measured.*

B. OpenQASM 3: The Need for Hybrid Quantum Programs

While OpenQASM 2 provided a great starting point, it has become apparent over time that some degree of classical logic and control flow, including conditionals and loops that might depend on the results of quantum measurements, is desirable for writing more complex quantum programs. For near-term applications, this allows to describe variational quantum algorithms, where the quantum circuit is part of a larger classical optimization loop. For long-term applications, classical feedback is essential for quantum error correction [25], [26]. As a consequence, OpenQASM has been extended to version 3 [21], which integrates classical logic and control flow into the IR.

Fundamentally, OpenQASM started as a quantum assembly language and step-by-step evolved into a more general-purpose quantum programming language by adding classical elements on top. It builds on its strong adoption in the quantum computing community and the corresponding compiler infrastructure previously built around OpenQASM 2. At the same time, the extension to OpenQASM 3 requires the implementation of traditional compiler optimizations on top of the IR, such as loop unrolling, constant propagation, or constant folding to generate efficient executable code for the target hardware. Given how ubiquitous OpenQASM is in the quantum computing community, this extension is a natural and sensible step to enable more complex quantum programs. However, it also requires the reimplementation of concepts that are already well-established and used for decades in classical compilers. This discrepancy has led to the development of QIR, which is discussed next.

C. QIR: Adding Quantum to the Classical

In contrast to OpenQASM, the *Quantum Intermediate Representation* (QIR, [22]) adopts a different strategy: It builds upon an established classical compilation infrastructure, augmenting it with quantum instructions. This approach allows QIR to utilize the existing classical compiler infrastructure, thereby inheriting its optimizations and transformations for classical code without additional effort. Specifically, QIR extends the classical IR used in the compiler framework *LLVM* [23]. It uses this IR to implement various compiler optimizations on this IR before, eventually, compiling it to target-specific assembly for execution. As such, LLVM IR can express arbitrary classical program logic, including functions, conditionals, and loops. However, it cannot express quantum computations by itself. To this end, QIR, as initially proposed by MICROSOFT, defines a set of additional functions that can be used to express quantum computations with the LLVM IR.

Example 2. *The relevant lines of a QIR program that represent the same circuit as in Ex. 1 are shown on the right in Fig. 1.*[1] *Initially, two arrays of two qubits and two classical bits are allocated. The pointers to those are stored in local variables on the stack labeled* `%q` *and* `%c`*, respectively. The Hadamard-gate is then applied to the first qubit. To this end, the first qubit is extracted from the qubit array and passed to the function* `@__quantum__qis__h__body`*. The CNOT-gate is applied in a similar fashion. Finally, the measurement of both qubits is performed.*

In order to speed up the adoption of QIR, multiple restrictions to QIR, so-called *profiles* have been defined that limit the expressiveness of QIR.[2] In its most restrictive form, the *base profile* only allows a sequence of quantum instructions that ends with the measurement of all qubits, which effectively makes it very similar to OpenQASM 2. The more permissive *adaptive profiles* allow the successive transition to fully support all features contained in LLVM IR. In its full generality, QIR is a proper superset of LLVM IR that builds on top of all the existing tooling and optimizations available for LLVM IR.

III. ADOPTING QIR: QUO VADIS?

The proposal of QIR is both logical and appealing: It makes sense to avoid reinventing the wheel and to build on decades of classical compilation expertise, especially, since quantum computations will inherently involve classical input and produce classical output. However, the adoption of QIR in the quantum computing software stack is a challenging task: What does it actually mean to adopt QIR in the stack? How do we transform these quantum-classical programs so that they can be executed? What about the actual execution? In the following, we discuss two fundamentally different directions of adopting QIR.

A. Parsing QIR Programs

In order for a quantum program to be executed, it must be transformed so that it complies with all the restrictions imposed by the hardware. Tools responsible for this transformation need to be able to accept the input quantum program and transform it accordingly [15]–[19]. When using QIR, tools may either receive the program as a text file containing QIR code or as an in-memory representation of the program, such as an *Abstract Syntax Tree* (AST). The former approach requires a parser that can turn the text-based representation into some in-memory representation. This in-memory representation can either be the QIR AST, in which case LLVM can be used for that job, or another custom/tool-specific IR. For a custom IR, the parser's (and, in fact, also the exporter's) complexity depends on the profile of QIR that should be supported.

Example 3. *For the base profile, it suffices to iterate over the lines to construct an in-memory representation of the resulting quantum circuit. For example, when parsing the QIR program from Ex. 2, the parser would need to track the assignment of variables (i. e.,* `%q`*,*`%0`*,*`%1`*,. . .) to their values to infer the*

---

[1] Note, that we chose to use modern LLVM syntax with opaque pointers here, as opposed to the syntax used in the QIR specification that is only compatible with LLVM versions prior to 16.

[2] https://github.com/qir-alliance/qir-spec/tree/main/specification/under_development/profiles (visited on 10/13/2024)

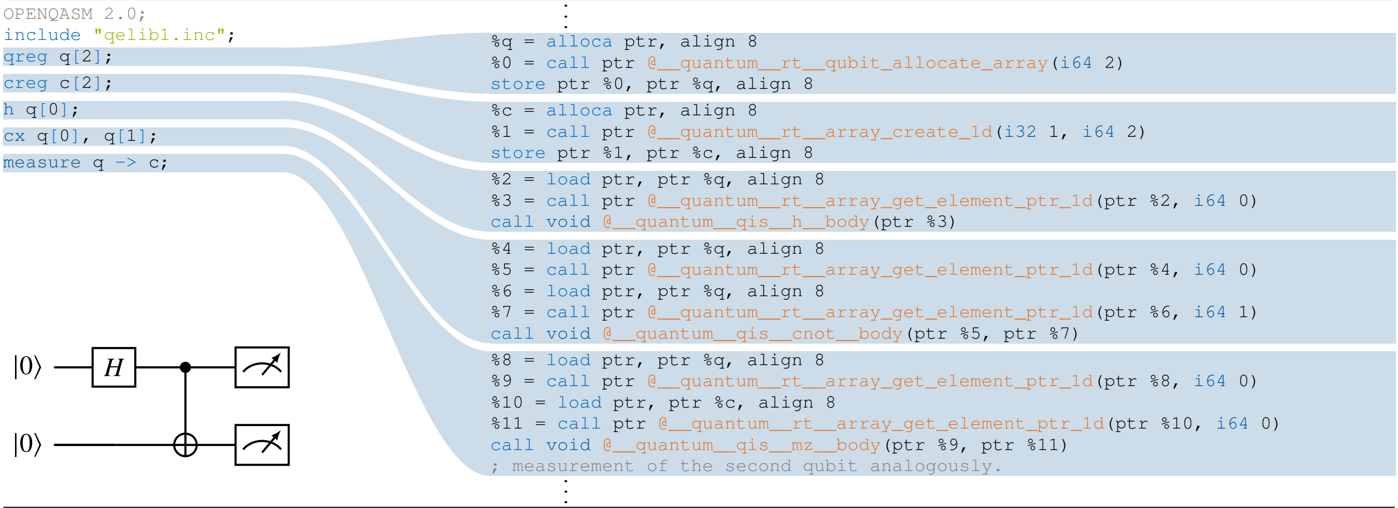

```
OPENQASM 2.0;
include "qelib1.inc";
qreg q[2];
creg c[2];
h q[0];
cx q[0], q[1];
measure q -> c;
```

```
      ⋮
%q = alloca ptr, align 8
%0 = call ptr @__quantum__rt__qubit_allocate_array(i64 2)
store ptr %0, ptr %q, align 8
%c = alloca ptr, align 8
%1 = call ptr @__quantum__rt__array_create_1d(i32 1, i64 2)
store ptr %1, ptr %c, align 8
%2 = load ptr, ptr %q, align 8
%3 = call ptr @__quantum__rt__array_get_element_ptr_1d(ptr %2, i64 0)
call void @__quantum__qis__h__body(ptr %3)
%4 = load ptr, ptr %q, align 8
%5 = call ptr @__quantum__rt__array_get_element_ptr_1d(ptr %4, i64 0)
%6 = load ptr, ptr %q, align 8
%7 = call ptr @__quantum__rt__array_get_element_ptr_1d(ptr %6, i64 1)
call void @__quantum__qis__cnot__body(ptr %5, ptr %7)
%8 = load ptr, ptr %q, align 8
%9 = call ptr @__quantum__rt__array_get_element_ptr_1d(ptr %8, i64 0)
%10 = load ptr, ptr %c, align 8
%11 = call ptr @__quantum__rt__array_get_element_ptr_1d(ptr %10, i64 0)
call void @__quantum__qis__mz__body(ptr %9, ptr %11)
; measurement of the second qubit analogously.
      ⋮
```

Figure 1. The quantum "Hello World" program, i. e., a circuit to create a Bell state (bottom left), expressed in OpenQASM 2.0 (top left) and QIR with dynamically allocated qubits (right). The corresponding lines of codes in each fragment are linked. For more details, see also Ex. 1 and 2.

respective qubit that is passed to a quantum instruction. The instructions themselves can be matched with a simple pattern and corresponding actions can be taken to gradually build up the quantum circuit.

On the one hand, the advantage of a custom parser is that one can avoid the dependency on LLVM, which should be a carefully considered decision because of the substantial size and complexity of LLVM. LLVM can significantly increase the build and maintenance overhead of the project. On the other hand, when not using LLVM, one has to reimplement all the optimizations and transformations that are already provided for LLVM IR "for free"——similar to the ongoing development of tools for OpenQASM 3. Additionally, when using a custom IR, one is limited to the capabilities of that existing IR. If that IR is not expressive enough to capture all the details of a quantum program, the QIR program cannot be fully translated and, hence, the tool cannot be used on these programs. As a consequence, the IR would have to add support for representing and handling all the concepts that QIR introduces, which might be a further significant effort in addition to reimplementing the optimizations and transformations.

B. Transforming QIR Programs

As a continuation of the first direction, after loading the quantum program, multiple passes in the software stack optimize and transform the program to meet all constraints imposed by the hardware. For the remainder of the paper, we assume that the program is represented as a QIR AST in memory. The support for QIR can, again, be realized in two different ways: Either the QIR is transformed directly, or the program is transpiled into another custom IR, transformed, and transpiled back into QIR. While the latter approach enables quick adoption of QIR, it carries the same deficits as parsing the text-based QIR file into a custom IR. For tools to support QIR directly, the AST of a QIR program must be handled and transformed. While this might be simple enough for the base profile, it becomes more complex for the adaptive profiles, which could have arbitrarily complex classical code in-between quantum instructions. This approach, however, has the advantage that the tool can directly take advantage of existing optimizations and transformations that are already implemented for QIR—the core motivation of an IR in a compiler.

Example 4. *One use case of classical* FOR*-loops is the application of a gate to multiple qubits. The following QIR snippet shows a simple* FOR*-loop that performs one Hadamard gate on the qubits* $0, \dots, 9$. *Since QIR builds on the LLVM infrastructure, it is straight forward to unroll any loops with statically known bounds in the QIR program. Hence, an optimization pass does not have to handle the* FOR*-loop, but sees only the ten individual Hadamard gates that are applied to the qubits.*

```
    %i = alloca i32, align 4
    store i32 0, ptr %i, align 4                    ; int i = 0
    br label for.header
for.header:
    %1 = load i32, ptr %i, align 4
    %cond = icmp slt i32 %1, 10                     ; i < 10
    br i1 %cond, label %body, label %exit
body:
    %2 = load i32, ptr %i, align 4
    call void @__quantum__qis__h__body(ptr inttoptr (i64 %2
      to ptr))                                     ; h(qi)
    %3 = load i32, ptr %i, align 4
    %4 = add nsw i32 %3, 1                          ; i++
    store i32 %4, ptr %i, align 4
    br label %for.header
exit: ...
```

While the previous sections discuss handling and transforming QIR programs themselves, the question remains how these transformed programs are eventually executed.

C. Executing QIR Programs

In contrast to the type of support for QIR discussed so far, the execution of QIR programs is fundamentally different. Another direction of adopting QIR enables the execution of QIR programs and can either be seen as a complement to the previously discussed options or as an orthogonal approach.

A file that contains LLVM IR bytecode can be executed directly with the `lli` tool provided by the LLVM project. This tool reads the bytecode and interprets it on the fly. Furthermore, the compiler `clang`, also part of the LLVM project, can compile the bytecode to machine code and emit an executable.

Obviously, when either of the tools is provided with a QIR file, they cannot handle the quantum instructions and will raise an error. However, this can be overcome by providing the missing definitions for the QIR extensions to LLVM. The resulting *quantum runtime* augments the classical LLVM runtime and, eventually, allows for the execution of arbitrary QIR programs.

Example 5. *One example that takes this approach is the* Catalyst Quantum Runtime *implemented by* XANADU *that incorporates their classical quantum circuit simulator* Lightning *[27], which is part of* PennyLane *[28]. Every function, such as* `@__quantum__qis__h__body`*, is implemented so that it modifies the internal state of the simulator to reflect the application of the respective gate. Note that these functions need not be implemented in LLVM IR, as this would be quite cumbersome. Instead, they can be provided as a C/C++/Rust implementation of the C interface functions, which is then compiled to LLVM just as any other regular classical progam code.*

This approach is orthogonal to the previous two options as it only concerns the implementation of quantum instructions, while the actual program structure is handled by the runtime. To this end, this approach is perfectly suited for integrating classical simulation techniques with QIR, especially ones developed in a language that compiles to LLVM such as C, C++, or Rust. The resulting binaries can be maximally optimized by the LLVM compiler infrastructure, which is a significant advantage.

IV. CHALLENGES IN COMPILING QIR PROGRAMS

The integration of a quantum IR like QIR into a compiler infrastructure is challenging and comes with a number of open research questions. Here, we detail two of those challenges.

A. Static and Dynamic Qubit Addresses

QIR provides the opaque pointer type `%Qubit*` for addressing individual qubits. In Ex. 2, the qubits were addressed dynamically by first allocating the respective qubits and retrieving dynamic qubit addresses that later can be used as arguments to quantum instructions. In contrast to dynamic addresses, there is also the possibility to address qubits statically.

Example 6. *When addressing the qubits statically instead of dynamically, the circuit from Fig. 1 can be rewritten in QIR in the following snippet. Especially, the lines for allocating the qubits disappear.*

```
       ⋮
call void @__quantum__qis__h__body(ptr null)
call void @__quantum__qis__cnot__body(ptr null, ptr
     inttoptr (i64 1 to ptr))
call void @__quantum__qis__mz__body(ptr null, ptr writeonly
     null)
call void @__quantum__qis__mz__body(ptr inttoptr (i64 1 to
     ptr), ptr writeonly inttoptr (i64 1 to ptr))
       ⋮
```

This type of addressing qubits can be useful when compiling a quantum program to an executable very close to hardware. In the end, the hardware only has a fixed number of qubits, and the compiler must ensure that the program does not exceed this number. So while a quantum program can be written with dynamic qubit addresses, the compiler must at some point assign the program's qubits to the hardware's qubits—a process very similar to register allocation in classical compilers. At the moment, this distinction between static and dynamic qubit addresses is not yet fully established in the QIR community, but it is an important aspect to consider when integrating QIR into a compiler infrastructure.

In the context of implementing a QIR runtime for a quantum circuit simulator, dynamic qubit addresses are the preferred way to address qubits. Most simulators support a variable number of qubits; the maximum feasible number might even depend on the input circuit itself, e. g., on the degree of entanglement (interdependence of several qubits' states) that is created during its execution. To support static qubit addresses, the runtime would either have to infer the number of qubits required for the simulation from the QIR program, such as via an attribute in the QIR file, or allocate qubits on the fly when it encounters a new qubit address that is not yet part of the simulated quantum state. Although this is not an inherent problem, supporting this feature in a simulator can still pose a significant challenge.

B. Hybrid Classical-Quantum Computing

As QIR is a superset of LLVM IR, arbitrary complex classical computations can be expressed in LLVM IR. Those heavy classical computations are preferably executed on dedicated classical hardware rather than on the classical co-processor of a quantum computer—if that is even feasible: The classical co-processor of quantum computers must be very fast, and special purpose hardware like FPGAs or ASICs are employed for those. These are incapable of executing arbitrary classical code, which is also not their purpose. Consequently, the question naturally arises for a hybrid classical-quantum program that contains quantum instructions as well as classical instructions how to decide which part of the code should be executed on the classical hardware and which part on the quantum hardware.

Especially, in the realm of error correction, where conditional gate applications based on intermediate measurements must be performed on the quantum computer to ensure low latency, the distinction is more complicated than just offloading only the quantum instructions. At the same time, it must be ensured, that the classical code offloaded to the quantum hardware can be executed in the required time frame to uphold the coherence of the qubits. Hence, as long as quantum computers cannot achieve arbitrary coherence of the qubits, there will always be programs that describe an infeasible execution and must be rejected.

V. CONCLUSIONS

QIR is one way to bridge the gap between purely classical and purely quantum computations. As an extension of LLVM IR, it builds on decades of compiler infrastructure and optimizations. However, the adoption of QIR in quantum computing software tools is a challenging task, as outlined in this article. Many open research questions need to be addressed. One framework for exploring solutions to these questions is the *Multi-Level Intermediate Representation* (MLIR, [29]), which is a natural choice for the next step in the evolution of QIR.

REFERENCES

[1] P.W. Shor, "Algorithms for quantum computation: Discrete logarithms and factoring," in *Symp. on Foundations of Computer Science*, IEEE Comput. Soc. Press, 1994. DOI: 10.1109/SFCS.1994.365700.

[2] Nikitas Stamatopoulos *et al.*, "Option pricing using quantum computers," *Quantum*, 2020. DOI: 10.22331/q-2020-07-06-291.

[3] Alberto Peruzzo *et al.*, "A variational eigenvalue solver on a photonic quantum processor," *Nature Communications*, 2014. DOI: 10.1038/ncomms5213.

[4] Stuart Harwood *et al.*, "Formulating and Solving Routing Problems on Quantum Computers," *IEEE Transactions on Quantum Engineering*, 2021. DOI: 10.1109/TQE.2021.3049230.

[5] Christa Zoufal, Aurélien Lucchi, and Stefan Woerner, "Quantum Generative Adversarial Networks for learning and loading random distributions," *npj Quantum Information*, 2019. DOI: 10.1038/s41534-019-0223-2.

[6] Lov K. Grover, "Quantum mechanics helps in searching for a needle in a haystack," *Phys. Rev. Lett.*, 2 1997. DOI: 10.1103/PhysRevLett.79.325.

[7] Ethan Bernstein and Umesh Vazirani, "Quantum complexity theory," *SIAM Jour. of Comp.*, 1997. DOI: 10.1137/S0097539796300921.

[8] Michel Boyer, Gilles Brassard, Peter Høyer, and Alain Tappa, "Tight bounds on quantum searching," in 2005. DOI: 10.1002/3527603093.ch10.

[9] Lieven M. K. Vandersypen *et al.*, "Experimental realization of Shor's quantum factoring algorithm using nuclear magnetic resonance," *Nature*, 2001. DOI: 10.1038/414883a.

[10] Dolev Bluvstein *et al.*, "Logical quantum processor based on reconfigurable atom arrays," *Nature*, 2023. DOI: 10.1038/s41586-023-06927-3.

[11] S. A. Moses *et al.*, "A race-track trapped-ion quantum processor," *Phys. Rev. X*, 4 2023. DOI: 10.1103/PhysRevX.13.041052.

[12] Hannah J. Manetsch *et al.*, *A tweezer array with 6100 highly coherent atomic qubits*, 2024. arXiv: 2403.12021.

[13] Martin Schulz, Laura Schulz, Martin Ruefenacht, and Robert Wille, "Towards the Munich Quantum Software Stack: Enabling Efficient Access and Tool Support for Quantum Computers," in *Int'l Conf. on Quantum Computing and Engineering*, IEEE, 2023. DOI: 10.1109/QCE57702.2023.10301.

[14] Lukas Burgholzer, Jorge Echavarria, Martin Schulz, Laura Schulz, and Robert Wille, "Building efficient software stacks for quantum computers: Experiences from the munich quantum software stack," in *Int'l Conf. on Quantum Computing and Engineering*, IEEE, 2024.

[15] Gushu Li, Yufei Ding, and Yuan Xie, "Tackling the Qubit Mapping Problem for NISQ-Era Quantum Devices," in *Int'l Conf. on Architectural Support for Programming Languages and Operating Systems*, ACM, 2019. DOI: 10.1145/3297858.3304023.

[16] Daniel Schoenberger, Stefan Hillmich, Matthias Brandl, and Robert Wille, "Using Boolean Satisfiability for Exact Shuttling in Trapped-Ion Quantum Computers," in *Asia and South Pacific Design Automation Conf.*, 2024. DOI: 10.1109/ASP-DAC58780.2024.10473902.

[17] Yannick Stade, Ludwig Schmid, Lukas Burgholzer, and Robert Wille, "An Abstract Model and Efficient Routing for Logical Entangling Gates on Zoned Neutral Atom Architectures," in *Int'l Conf. on Quantum Computing and Engineering*, arXiv, 2024. arXiv: 2405.08068.

[18] Deborah Volpe, Nils Quetschlich, Mariagrazia Graziano, Giovanna Turvani, and Robert Wille, "A Predictive Approach for Selecting the Best Quantum Solver for an Optimization Problem," in *Int'l Conf. on Quantum Computing and Engineering*, arXiv, 2024. arXiv: 2408.03613.

[19] Robert Wille *et al.*, "The MQT Handbook: A Summary of Design Automation Tools and Software for Quantum Computing," in *Int'l Conf. on Quantum Software*, 2024. arXiv: 2405.17543, A live version of this document is available at https://mqt.readthedocs.io.

[20] Andrew W. Cross, Lev S. Bishop, John A. Smolin, and Jay M. Gambetta, *Open Quantum Assembly Language*, 2017. arXiv: 1707.03429.

[21] Andrew W. Cross *et al.*, "OpenQASM 3: A broader and deeper quantum assembly language," *ACM Transactions on Quantum Computing*, 2022. DOI: 10.1145/3505636.

[22] *Qir specification*, version 0.1, QIR Alliance: https://qir-alliance.org, 2021.

[23] C. Lattner and V. Adve, "Llvm: A compilation framework for lifelong program analysis & transformation," in *Int'l Symp. on Code Generation and Optimization*, IEEE Computer Society, 2004. DOI: 10.1109/CGO.2004.1281665.

[24] Michael A. Nielsen and Isaac L. Chuang, *Quantum Computation and Quantum Information*, 10th anniversary ed. Cambridge University Press, 2010.

[25] Earl T. Campbell, Barbara M. Terhal, and Christophe Vuillot, "Roads towards fault-tolerant universal quantum computation," *Nature*, 2017. DOI: 10.1038/nature23460.

[26] Tom Peham, Ludwig Schmid, Lucas Berent, Markus Müller, and Robert Wille, *Automated Synthesis of Fault-Tolerant State Preparation Circuits for Quantum Error Correction Codes*, 2024. arXiv: 2408.11894.

[27] Ali Asadi *et al.*, *Hybrid quantum programming with PennyLane Lightning on HPC platforms*, 2024. arXiv: 2403.02512.

[28] Ville Bergholm *et al.*, *Pennylane: Automatic differentiation of hybrid quantum-classical computations*, 2022. arXiv: 1811.04968.

[29] Chris Lattner *et al.*, "MLIR: Scaling compiler infrastructure for domain specific computation," in *Int'l Symp. on Code Generation and Optimization*, 2021. DOI: 10.1109/CGO51591.2021.9370308.