Why Are We Unrolling? The Importance of Structured Quantum Programs for Compilation

Damian Rovara*†, Daniel Haag*‖, Mark Koch‡§, Josh Izaac¶, Seyon Sivarajah§, Robert Wille†‖, Agustín Borgna§, Lukas Burgholzer†‖, Brad Chase**, Olivia Di Matteo††, David Ittah¶
*Technical University of Munich, Munich, Germany
† Corresponding author: damian.rovara@tum.de
‡University of Oxford, Oxford, UK
§Quantinuum, Cambridge, United Kingdom
¶Xanadu, Toronto, ON, M5G 2C8, Canada
‖MQSC, Garching near Munich, Germany
**Unitary Foundation, San Francisco, CA, USA
††The University of British Columbia, Vancouver BC, Canada

***Abstract*—As quantum software stacks scale up to support future fault-tolerant quantum hardware and algorithms, quantum compilation is becoming an increasingly important component of the stack. How do we ensure that our software stacks support dynamic algorithms—including patterns such as mid-circuit measurement feedforward and repeat-until-success—with hundreds of logical qubits and billions of quantum operations? To do so, we must re-think how we represent quantum programs beyond straight-line circuits, to representations that include classical structure and dynamism, and make this the default representation to consider when performing quantum compilation. In this work, we present important patterns and algorithms from fault-tolerant quantum applications which admit a structured representation that we argue is crucial to preserve, and set a challenge to the community to compile such representations without unrolling them into straight-line quantum circuits. We also explore the status quo of structured program support in quantum software, and ask ourselves the rhetorical question: how much more efficient can we make quantum compilation tooling when we take into account the additional information from classical structure?**

***Index Terms*—quantum compilation, quantum software, structured quantum programs, benchmarking**

I. Introduction

The accelerating pace of quantum hardware towards fault-tolerance necessitates software capable of orchestrating the execution of quantum algorithms at scale. Resource estimates for solving real-world problems with active error correction often yield physical qubit counts on the order of millions and gate counts exceeding billions [1]–[3]. Managing such a large amount of operations is a complex scheduling problem; determining that sequence of operations in the first place is a separate challenge addressed by *quantum compilation*.

Quantum compilation translates an implementation of a quantum algorithm into executable instructions. Effective quantum compilers enable users to implement algorithms with minimal concern for execution details such as error correction, gates on individual qubits, or processor architecture. Historically, quantum algorithms were primarily expressed as explicit logical circuits. Thus, "quantum compilation" was synonymous with quantum circuit synthesis and optimization, which unrolls a fully-specified circuit into a flat list of gates (or graph) and applies a sequence of transformations. These transformations are usually equivalence-preserving and may include circuit synthesis and decomposition, rewriting and pattern-matching optimization, transpilation into different gate sets, and modifications based on hardware-specific constraints.

Significant recent effort has been invested in programming frameworks that raise the abstraction level beyond the circuit model [4]–[11], allowing algorithms to be expressed using familiar constructs like subroutines, loops, and classical control flow. This raises the question of how to compile and optimize such programs. Hardware execution involves streaming a list of instructions to a device. The list can be generated in advance, by unrolling loops or creating multiple program versions to account for branching, and optimized with circuit transformations. However, this has scalability implications as many circuit transformations require solving computationally hard problems (e.g., qubit allocation is NP-complete [12] and qubit routing is NP-hard [13]); incorporating error correction adds yet another layer of complexity. Moreover, any change in program parameters incurs recompilation.

Parametric compilation is one approach to avoid recompilation [14]. In parametric compilation, programs are compiled with symbolic input parameters, whose values are provided at execution time — avoiding the need for re-compilation with new parameter values. Python-based frameworks such as JAX [15], Catalyst [16], and Qrisp [6] often implement parametric compilation via what's known as just-in-time (JIT) compilation. In these frameworks, the first time a program is executed in the Python interpreter, tracers are passed through a program instead of concrete input values, which then record the program's instruction stream (possibly including structure). This representation is compiled and cached for reuse in subsequent executions, avoiding recompilation provided the type signature of the inputs doesn't change.

However, avoiding recompilation is not the only challenge. Some quantum algorithms fundamentally *cannot* be expressed

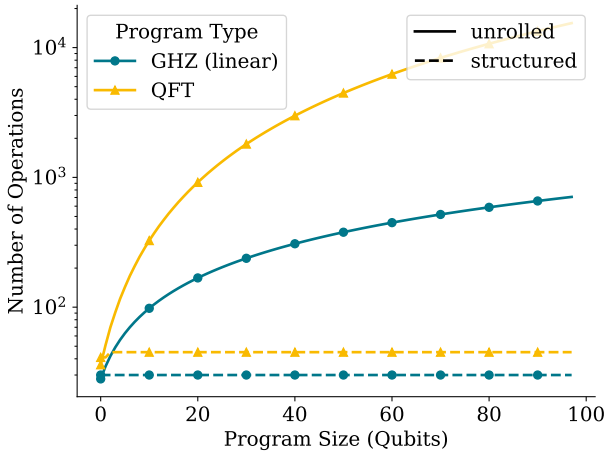

Fig. 1: Preserving program structure can enable constant compilation times, even when the problem instance size increases. On the other hand, unrolling loops generally leads to polynomial or even exponentially-increasing program size. Programs used for this comparison were generated using `mqt-cc` [25]. See section IV for details on benchmark programs.

as one flat circuit. This includes: repeat-until-success patterns in circuit synthesis [17], block encodings for Hamiltonian simulation [18], [19] and quantum singular value transformations [20]; magic state distillation [21]; active error correction [22]; and Shor's algorithm [23]. Some intermediate representations, enabled by classical compiler infrastructure, can preserve and exploit structure in these algorithms [6], [8], [16], [24]–[28], leading to compilation times and program sizes independent of problem size. Such constant scaling, depicted graphically for sample problems in Figure 1, was demonstrated empirically in multiple frameworks [6], [29], [30].

A major roadblock is that, unless unrolled, programs with structure are incompatible with compilers that only support flat circuit transformations. Novel, dedicated tooling is required to both support and exploit such structure during optimization. While quantum circuit optimization is well-studied, there are fewer approaches to optimization of structured quantum programs. Often, such optimizations are the result of human intuition and manual implementation [30]. Existing systematic methods rely on underlying mathematical structures that limit their utility. For example, assertion-based optimization uses the Z3 theorem prover to establish pre- and post-conditions that enable removal of trivial gates and subroutines [31]. Relational analysis of program input and output variables can be used to formulate techniques such as phase folding and perform optimizations around loops, but encounters challenges with gates that create superpositions [32].

The tension between preserving program structure and flattening it into a linear instruction stream is not unique to quantum computing; it has shaped classical accelerator compilation for decades. Polyhedral compilers such as *Pluto* and *PPCG* [33], [34] retain loop nests as an explicit iteration space instead of unrolling them. The same principle applies to intermediate representations (IRs) built specifically for accelerator generation: recent designs deliberately keep loops and conditionals in the IR rather than eliminating them up front to enable further optimizations—lowering control flow only as a final step [35], [36]. Database query engines are a further classical instance: a query is optimized as a relational-algebra plan and only then compiled to machine code, so that the structural information in the plan is exploited before any per-tuple work is materialized [37]. In contrast, reconstructing loops from already flattened code is an active field of research, aiming to reduce code size and recover analysis precision [38], [39]. Where flattening does occur, it is a local and deliberate choice—unrolling a small inner loop for throughput—not a generic expansion of the program [40].

In light of today's landscape, we argue that preserving program structure is not just optional, but *essential* for scalable quantum compilation. This necessitates research on both the representation and optimization of structured programs. In this work, we make three contributions towards that goal:

- We provide a formal definition of *structured quantum programs* and their constituent elements.
- We survey the software landscape and assess current frameworks' level of support for structured programs
- We propose a diverse set of programs, with varying amounts of structural elements, to serve as a benchmark set for current and future compilers. We highlight explicit examples where preserving structure in the representation is essential for compilation.

II. Structured Quantum Programs

A. Definition

A *quantum program* is a computer program wherein one or more instructions are executed on a quantum device. A characteristic feature of *structured* quantum programs is a representation with hierarchical structure beyond a flat list of instructions, such as subroutines or structured control flow. More concretely, structured program execution does not proceed in a strictly sequential manner, as formalized below.

Definition 1. *Let P be a quantum program, and c_P an integer-valued program counter indicating the line of P currently being executed. P is a* structured quantum program *if c_P does not evolve exclusively by unit increments, i.e., there exist execution steps where $c_P \not\mapsto c_P + 1$.*

Seven key primitives of structured programs are outlined in Table I, and examples of equivalent structured and non-structured programs are presented in Table II. Program types are further subdivided based on data dependence. If the quantum instructions in P are fully determined before execution, P is a *static* quantum program. In *dynamic* quantum programs, c_P, instructions at c_P, or instruction parameters depend on classical values or quantum measurements resolved at runtime.

Many primitives from Table I occur naturally in quantum algorithms, though their use may be optional. A concrete example is included below in a pseudo-Pythonic implementation

TABLE I: Elements of structured quantum programs. The Tier, described in subsection III-A, differentiates between elements that are merely advantageous (Tier 1) vs. required (Tier 2) for a complete, universal program representation. Note that conditioning and iteration can appear together in a single program construct, like a while loop.

Tier	Primitives of structured programs	Description
1	StatIter	Program uses iteration (loops) with a constant number of repetitions.
2	DynIter	Iteration (loop) termination conditions depend on values computed at runtime
1	DynVal	Gates use classical values computed at runtime (e.g., rotation angles from arrays)
1	CCond	Conditional depends on classical values that are not results of quantum measurements
2	QCond	Conditionals depend on results of one or more quantum measurements
1	DynAlloc	Qubits are allocated at runtime (e.g., inside loop bodies)
2	DynIdx	Gates are applied to qubits with non-constant indices (e.g., loop variable)

TABLE II: Examples of static and dynamic (un)structured quantum programs using OpenQASM 2 [41] and OpenQASM 3 [8] syntax. In both unstructured programs, the program counter increases linearly. In the structured static program, the counter jumps back when the `ghz_cnots` subroutine executes, then follows loop execution. In the structured dynamic program, the counter cycles through the `while` loop; note such dynamically-bounded `while` loops are not supported in OpenQASM 2.

Static

Unstructured

```
qreg q[4];
h q[0];
cx q[0], q[1];
cx q[0], q[2];
cx q[0], q[3];
```

Structured

```
def ghz_cnots(q) {
    for i in [1:3] {
        cx q[0], q[i];
    }
}
qreg q[4];
h q[0];
ghz_cnots(q);
```

Dynamic

Unstructured

```
qreg q[2];
creg c[1];
ry(0.1) q[0];
measure q[0] -> c;
if (c == 1) x q[1];
```

Structured

```
qreg q[2];
creg c[1];
while (c != 1) {
    ry(0.1) q[0];
    measure q[0] -> c;
    reset q[0];
}
x q[1];
```

of Shor's algorithm, based on the implementation in Ref. [30]. Instances of each primitive are noted in the comments.

```
def shors_algorithm(N):
  a = randint(2, N - 1)

  # CCond: conditional on dynamic classical value
  p = gcd(a, N)
  if p != 1:
    return p, N // p

  # DynAlloc: qubit register size based on bitwidth
  n_bits = floor(log2(N)) + 1
  est_q = qalloc(1)
  target_q = qalloc(n_bits)
  aux_q = qalloc(n_bits + 2)

  meas = zeros((n_bits, ))
  cumul_phase = 0.0
  idx = 0
  pow_a = a

  def run_qpe():
    PauliX(target_q[-1])

    # DynIter,CCond: loop termination based on
    n_bits, a
    while idx < n_bits and pow_a != 1:
      Hadamard(est_q)

      # DynIdx: qubit indices in loops based on
    n_bits
      # DynVal: gates will depend on binary form of
    a
      ctrl_ua(N, pow_a, est_q, target_q, aux_q)

      PhaseShift(cumul_phase, est_q)
      Hadamard(est_q)

      meas[idx] = measure(est_q, reset=True)

      # QCond: conditional on measurement result
      if meas[idx] == 1:
        cumul_phase += -2 * pi / 2 ** (idx + 1)

      idx += 1
      pow_a = a ** (2**idx) % N
    return meas

  p, q = 0, 0

  # StatIter: statically-bounded loop
  for _ in range(100):
    sample = run_qpe()
    guess_r = samples_to_order(sample, N)

    if guess_r % 2 == 0:
        # ...

  return p, q
```

This is a dynamic structured quantum program: a random value `a`, selected at runtime, affects classical conditionals, loop termination conditions, and quantum gates. Array and qubit register sizes depend on the bit width of N, and there are conditionals based on measurement results. Quantum gates are executed in multiple program regions. Most are in the `while` region's modular exponentiation subroutine (`ctrl_ua`), wherein gates, loops, and register indexing are based on `a` and N. Optimizations using the runtime value of `a` can also be incorporated (e.g., removing controlled operations based on its binary representation, as done in [30]).

In Shor's algorithm, the circuit depends intimately on a and N so it will be different each time the algorithm executes. Preserving structure where possible allows compilation of a more compact representation, which is more scalable

than immediately unrolling. Of course, some algorithms may eventually be unrolled prior to hardware execution. However, the key is that sometimes structural elements are essential: the dynamic loop boundary (line 24) means the algorithm *cannot* be represented by a flat circuit before execution. In what follows, we provide a more comprehensive exploration of how the structural elements in Table I impact compilation.

III. Compiling with Structure: Goals and Requirements

This section discusses the core aspects of compiling structured quantum programs, dividing structural elements into a system of two tiers and discussing their requirements.

A. A Taxonomy of Structured Quantum Program Features

To address the two core concerns of computational bottlenecks and expressivity, we revisit the structural elements defined in Table I. Rather than grouping elements by syntactic similarity, we categorize them by their impact on the quantum compilation pipeline. In doing so, we create a taxonomy of structural elements consisting of a baseline and two distinct tiers: those enabling more efficient compilation, and those strictly required for universal execution of input programs.

Tier 0: Straight-Line Quantum Circuits

As a baseline, we define Tier 0 features to represent standard quantum circuits devoid of high-level control flow, i.e., purely sequential application of quantum gates and measurements. While this representation is sufficient for many near-term algorithms, it fails to capture the complexity and scalability required by fault-tolerant quantum circuits or advanced algorithms without resorting to massive, unrolled circuits.

Tier 1: Structure as an Enabler

This tier of structural elements comprises features that are not a strict necessity for purely expressing a quantum program, but provide significant advantages by enabling advanced optimizations, reducing code size, and encoding detailed semantic intent. The defining aspect of this tier is that its structural elements can, in principle, be deterministically translated into straight-line circuit primitives at compile time (e.g., by fully unrolling statically-bounded loops). Consequently, existing quantum circuit compilers can achieve a baseline level of compatibility simply by implementing such translations as preprocessing steps. However, to truly benefit from this tier, compilers must evolve beyond simple unrolling and preserve these elements natively in their IRs. This tier includes the following elements and their benefits:

- *Statically-bounded loops* reduce redundancies in the IR.
- *Conditionals on classical values* allow the conditional application of quantum operations in a way that can be analyzed statically.
- *Dynamic classical values* allow parametrized operations to be given in a more general representation.
- *Dynamic qubit allocation* allows qubits to be confined to the parts of the code where they are required.

TABLE III: Comparison of common compilation frameworks and their ability to handle structural primitives.

Framework	StatIter	DynIter	Dyn Val	CCond	QCond	Dyn Alloc	Dyn Idx
Qiskit	✓	✗	✓	✓	✓	✓	✗
TKETv1	✗	✗	✗	✓	✓	✗	✗
TKETv2	✓✓	✓✓	✓✓	✓✓	✓✓	✓	✓✓
Catalyst	✓✓	✓✓	✓✓	✓✓	✓✓	✓✓	✓✓
mqt-cc	✓	✓	✓✓	✓✓	✓✓	✓✓	✓

✗: incompatible ✓: compatible ✓✓: exploitable

Tier 2: Structure as a Requirement

In contrast to the enabling nature of Tier 1, this tier consists of structural elements that are strictly required for the universal execution of complex quantum programs. The defining characteristic of these elements is that they inherently depend on dynamic or runtime information. Consequently, they cannot be preprocessed or unrolled into straight-line circuits under any circumstances. Compilers must explicitly support and handle these features to even be compatible with the programs that employ them. This tier includes the following elements, as well as the reasons why they cannot be reduced to straight-line circuits:

- *Dynamically-bounded loops* cannot be unrolled at compile time as iteration counts depend on runtime values and cannot be computed in advance.
- *Dynamic qubit indexing* obscures data dependencies and cannot be replaced by static qubit indexing.
- *Conditionals on measurement results* limit static analysis of conditional branches and loops and cannot be expressed by static circuit primitives.

B. A Gap in Representation

To prepare for the rising complexity of quantum programs, a quantum compilation framework must be able to address both tiers of structural elements presented above. Unfortunately, the current landscape of quantum software is heavily fragmented, with varying levels of support for these critical elements. Modern quantum execution relies on a multi-stage compilation pipeline, ranging from high-level front-end programming languages down through several IRs to formulate the target instructions for the quantum hardware. In this pipeline, the actual support for structured quantum programs is ultimately dictated by the "lowest common denominator" among all involved layers. It is simply not enough for only certain stages or intermediate formats to be capable of expressing these features. To fully exploit structured programs, these elements need to be broadly supported across the entire toolchain. More importantly, compilers must actively be able to transform their IR while preserving these features so that optimization passes can leverage them directly.

To assess the current state of this software pipeline, Table III evaluates the support for structural primitives across a variety of quantum compilation frameworks. A single checkmark indicates *compatibility*, i.e., that a framework can consume programs containing the primitive and preserve runtime semantics, even if structure is lost in the process. A double

checkmark indicates that structure is preserved and the compiler is currently engineered to actively transform and optimize these primitives, *exploiting* the embedded information to produce more efficient executables. While the table shows a positive trend of frameworks beginning to explicitly tackle structured program compilation, the data points clearly show that support remains highly fragmented and incomplete.

An examination of popular formats, such as QIR [24] and OpenQASM 3 [8], reveals a primary cause for this discrepancy. Fundamentally, neither QIR nor OpenQASM 3 are designed purely as intermediate representations for aggressive compiler transformations; rather, they serve primarily as exchange formats between the compiler stack and the quantum hardware. Because quantum hardware is historically very restricted in the types of operations it supports natively, wide compatibility relies heavily on straight-line circuits. Consequently, these formats often err on the side of caution. They heavily isolate structured elements, for instance by limiting compilers to strictly flat profiles like the QIR Base Profile, or by explicitly marking features like dynamic qubit indexing as optional in OpenQASM 3. Since the ultimate target architectures often lack native support for structured primitives, there has historically been little incentive for compiler developers to support them. Due to the "lowest common denominator" effect, many toolchains based on these standard exchange formats simply bypass structured control flow altogether.

However, a shift is occurring with IRs that are more deeply rooted in classical compiler infrastructure. Frameworks building upon MLIR [16], [25], [27], as well as quantum-focused IR frameworks like HUGR [42] and Kirin [43], are designed on top of mature classical compilation workflows or take inspiration from them. As a result, they more naturally incorporate and retain support for structured primitives in optimization passes, pointing toward a more robust handling of advanced control flow natively within the compiler stack.

Addressing the compilation of structured quantum programs therefore boils down to two key aspects: The preservation of structure and its exploitation to produce more efficient executables. To drive this evolution, it is critical to rely on comprehensive benchmarks. Benchmarks are the fundamental tool needed to continuously evaluate the state of the art, identify the most pressing gaps in current support, and map the road ahead for fully structured quantum compilation.

IV. A Challenge for Structured Quantum Program Compilation

To move the ecosystem beyond its current reliance on straight-line circuits, we propose an open challenge to the developers of quantum compilers. This challenge introduces a comprehensive benchmark suite of quantum programs that employ varying degrees of structural elements. We align the evaluation of compilation frameworks with two core requirements: *support* and *utilization*. First, to establish baseline *support* within state-of-the-art compiler frameworks, representations must be compatible with the individual elements and compilers must handle them effectively. Second, to drive the *utilization* of these concepts, compilers must be able to natively leverage these elements when compiling structured programs.

Table IV presents the complete benchmark suite, detailing the structural elements required by each program. These programs were selected to represent a wide range of realistic algorithms covering the different structural elements discussed above.

Some programs mainly feature statically-bounded loops (StatIter), conditionals on classical values (CCond), and other static elements that can be translated to straight-line circuits during preprocessing. While this makes baseline *support* for these features trivial, compilers must specifically preserve and leverage the provided structure to truly drive *utilization*.

For instance, *Grover's Search* relies heavily on *statically-bounded loops* (StatIter). The number of iterations can be precomputed, allowing the program to be represented as a fixed number of N iterations over the *oracle* and *diffusion* steps typically employed for the algorithm. To achieve baseline *support*, a naive quantum compiler can handle this by transforming the program into a flat sequence of N consecutive repetitions. However, doing so forces the optimizer to independently process all gates each iteration, a process that grows exponentially inefficient since the number of iterations scales as $N = \mathcal{O}(\sqrt{2^n})$ for an n-qubit search space. Instead, by preserving the structure and performing optimizations only within the loop body, an advanced compiler can *utilize* the structure to drastically reduce memory and runtime overhead.

Conversely, other benchmarks feature dynamic elements that intrinsically rely on (quantum) runtime results, making it impossible to flatten them into straight-line circuits. To sufficiently *support* these algorithms, compilers must explicitly implement native handling for elements such as dynamically-bounded loops (DynIter), conditions on measurements (QCond), or dynamic qubit indexing (DynIdx).

Program classes such as *Repeat-Until-Success* provide a classic baseline for testing these dynamic capabilities. The contained loops require dynamically-bounded execution (DynIter) to validate that a desired state has been successfully prepared, an operation has been correctly applied, or a specific parity check has passed. Explicit representations and targeted optimizations are required to correctly parse and process these program classes. Importantly, successfully utilizing these elements is not solely a compiler problem. Such advanced control-flow operations require close coordination with the actual control systems of the target quantum hardware [44]. A compiler can only represent and optimize dynamic operations to the extent that the underlying physical devices and their classical control logic can efficiently execute them at runtime.

A benchmark suite is more than a collection of examples for individual structural features; its true value lies in allowing developers to evaluate the scalability of their frameworks. To facilitate comprehensive scalability studies, many program classes included in this challenge can be scaled to an arbitrary size. *GHZ preparation*, for instance, can be defined over any number of qubits. By instantiating these challenge programs at varying sizes, compiler developers can stress-test the compila-

TABLE IV: List of benchmark programs and their structural features.

Algorithm	StatIter	DynIter	DynVal	CCond	QCond	DynAlloc	DynIdx	Arb. Size
Basic Subroutines & Primitives								
GHZ (linear & star)	✓						✓	✓
Q. Multiplexers	✓		✓	✓			✓	✓
Block Encoding					✓			
Toffoli-heavy Circuits	✓					✓	✓	✓
Parallelization with quantum fan-out	✓		□			✓	✓	✓
Teleportation					✓			
Transforms & Arithmetic								
QFT	✓		□				✓	✓
QPE	✓		✓				✓	✓
Iterative QFT	✓		□				✓	✓
Iterative QPE	✓		✓				✓	✓
QFT adder (quantum input, two registers)	✓		□				✓	✓
QFT adder (classical input, single register)	✓		✓	□			✓	✓
Controlled multiplication modulo N	✓	✓	✓	□			✓	✓
Heuristics, Optimization & Search								
Grover's Search	✓			□				✓
Grover's Search with Weak Measurement		✓		□	✓			✓
VQE Ansatz with Fixed Repetitions	✓		✓					✓
QAOA with Fixed Repetitions	✓		✓					
VQE	✓		✓		✓			✓
ML-QAE		✓	✓	✓	✓		✓	
Complex Algorithms & Simulation								
Shor's Algorithm	✓	□	✓				✓	✓
X-Ray Absorption Spectroscopy	✓		✓	✓			✓	
Error Correction & Fault Tolerance								
Repeat-Until-Success		✓			✓			
Magic State Distillation	✓	✓			✓		✓	✓
Logical State Preparation	✓	✓			✓		✓	✓
Syndrome Measurement and Correction	✓	✓			✓		✓	✓
Measurement-based quantum computation	□		✓		✓		□	

✓: Element required Blank: Element not required □: Implementation-dependent

tion times and memory footprints of their pipelines, ensuring their tools remain robust even for massive program instances.

The complete benchmark suite is distributed using the `jeff` exchange format [26]. `jeff` has been specifically created to exchange structured quantum programs between different compilers and frameworks while ensuring native support for all proposed structural elements. This makes it as easy as possible for developers to get started with the benchmark suite by translating from one common standard to their format of choice. Conversions from `jeff` to major open-source SDKs are either already available or currently under development, including PennyLane [5] / Catalyst [16], HUGR [42], Kirin [43], and the MQT Compiler Collection (`mqt-cc`) [25].

V. Evaluation: case study

To trace the compatibility of existing compilation frameworks with structured quantum programs, we investigate to what degree the *Repeat-Until-Success* (RUS) benchmark can be processed by state-of-the-art compilers. To this end, we divide the compilation pipeline into three distinct steps: *Parsing* programs containing structural elements, *optimizing* them to reduce total resource requirements, and *lowering* them to a format that can be executed on physical quantum hardware.

By investigating the individual compiler API specifications, we come to the following conclusions: *Qiskit's* [4] API does not allow the specification of dynamically-bounded `while` loops. Consequently, for *Qiskit*, the compilation pipeline of RUS fails already at the *parsing* level. `mqt-cc`, on the other hand, explicitly allows the specification of structured programs with dynamically-bounded loops and preserves them during optimization. Whether the lowering step succeeds depends on the targeted output format: when targeting the *QIR Adaptive Profile*, `mqt-cc` can lower such structured programs to an executable format, whereas the *QIR Base Profile* does not support these constructs and the pipeline fails at the lowering level. For `mqt-cc`, the compilation pipeline of RUS is therefore complete in principle, but only for a subset of the supported output formats Going further, the *TKETv2* compiler can compile structured programs in the Guppy language [11] to HUGR, and lower to *streaming execution* on the *Helios* quantum computer [44], allowing RUS to follow the full compilation pipeline down to the execution on quantum hardware. Similarly, the Catalyst compiler supports such structure at the language level (via PennyLane), compiler representation and transformations, and execution level — supporting an array of simulators but no current hardware devices.

This case study shows that, while some individual compilation frameworks already fully or partially support structured quantum programs, compatibility is far from universal and further effort must be invested in state-of-the-art compilers to make the execution of structured quantum programs possible.

VI. Conclusion and next steps

In this work, we have motivated the transition from simple straight-line quantum circuits to *structured quantum programs*. We have defined the individual aspects of structured programs and constructed a taxonomy to facilitate the development of next-generation compilers that are compatible with these more advanced structures. Incorporating structure into quantum compilers increases their efficiency and allows quantum computing users to develop more scalable and readable programs.

Developing more general, automated higher-level optimization strategies that leverage program structure is an exciting and open area of research. The challenge programs we provide [26] are a call to the community to start incorporating support for structured programs in their quantum compilers, so that future quantum programs can natively exploit these features, much like classical structured control flow has long been a staple in traditional software development.

Acknowledgments

D.R., D.H., R.W., and L.B. 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 ("MILLENIONSGA1" EU Project), and the Munich Quantum Valley, which is supported by the Bavarian state government with funds from the Hightech Agenda Bayern Plus. Furthermore, this work was supported by the BMFTR under grant numbers 13N17298 (SYNQ) and 01MQ25001I (FullStaQD), the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under grant numbers 563402549 and 563436708, and the Austrian Research Promotion Agency (FFG) together with the states of Upper Austria and Tyrol within the COMET module Quantum Algorithm Engineering (FFG) under grant number 923923.

ODM acknowledges funding from NSERC, the Canada Research Chairs program, and UBC.

References

[1] M. E. Beverland, P. Murali, M. Troyer *et al.*, "Assessing requirements to scale to practical quantum advantage," 2022. [Online]. Available: https://arxiv.org/abs/2211.07629

[2] A. A. Agrawal, J. Job, T. L. Wilson *et al.*, "Quantifying fault tolerant simulation of strongly correlated systems using the fermi-hubbard model," 2024. [Online]. Available: https://arxiv.org/abs/2406.06511

[3] N. Nguyen, T. W. Watts, B. Link *et al.*, "Quantum computing for corrosion-resistant materials and anti-corrosive coatings design," 2025. [Online]. Available: https://arxiv.org/abs/2406.18759

[4] A. Javadi-Abhari, M. Treinish, K. Krsulich *et al.*, "Quantum computing with Qiskit," 2024.

[5] V. Bergholm, J. Izaac, M. Schuld *et al.*, "Pennylane: Automatic differentiation of hybrid quantum-classical computations," *arXiv preprint arXiv:1811.04968*, 2018.

[6] R. Seidel, S. Bock, R. Zander, M. Petrič, N. Steinmann, N. Tcholtchev, and M. Hauswirth, "Qrisp: A framework for compilable high-level programming of gate-based quantum computers," 2024. [Online]. Available: https://arxiv.org/abs/2406.14792

[7] C. Developers, *Cirq*. Zenodo, Apr. 2025. [Online]. Available: https://zenodo.org/doi/10.5281/zenodo.4062499

[8] A. Cross, A. Javadi-Abhari, T. Alexander, and Dothers, "OpenQASM 3: A Broader and Deeper Quantum Assembly Language," *ACM Transactions on Quantum Computing*, 2022.

[9] K. Svore, A. Geller, M. Troyer *et al.*, "Q#: Enabling scalable quantum computing and development with a high-level dsl," in *Proceedings of the Real World Domain Specific Languages Workshop 2018*, 2018.

[10] B. Bichsel, M. Baader, T. Gehr, and M. Vechev, "Silq: a high-level quantum language with safe uncomputation and intuitive semantics," in *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2020, p. 286–300.

[11] M. Koch, A. Lawrence, K. Singhal, S. Sivarajah, and R. Duncan, "GUPPY: Pythonic Quantum-Classical Programming," 2025.

[12] M. Y. Siraichi, V. F. d. Santos, C. Collange, and F. M. Q. Pereira, "Qubit allocation," in *Proceedings of the 2018 International Symposium on Code Generation and Optimization*, 2018.

[13] T. Ito, N. Kakimura, N. Kamiyama, Y. Kobayashi, and Y. Okamoto, "Algorithmic theory of qubit routing in the linear nearest neighbor architectures," *ACM Transactions on Quantum Computing*, vol. 6, no. 3, May 2025. [Online]. Available: https://doi.org/10.1145/3722119

[14] P. J. Karalekas, N. A. Tezak, E. C. Peterson, C. A. Ryan, M. P. da Silva, and R. S. Smith, "A quantum-classical cloud platform optimized for variational hybrid algorithms," *Quantum Science and Technology*, 2020.

[15] J. Bradbury, R. Frostig, P. Hawkins *et al.* (2018) JAX: composable transformations of Python+NumPy programs. [Online]. Available: http://github.com/jax-ml/jax

[16] D. Ittah, A. Asadi, E. O. Lopez, S. Mironov, S. Banning, R. Moyard, M. J. Peng, and J. Izaac, "Catalyst: a Python JIT compiler for auto-differentiable hybrid quantum programs," *Journal of Open Source Software*, vol. 9, no. 99, p. 6720, 2024. [Online]. Available: https://doi.org/10.21105/joss.06720

[17] A. Paetznick and K. M. Svore, "Repeat-until-success: non-deterministic decomposition of single-qubit unitaries," *Quantum Inf. Comput.*, vol. 14, pp. 1277–1301, 2013. [Online]. Available: https://api.semanticscholar.org/CorpusID:15282418

[18] G. H. Low and I. L. Chuang, "Optimal hamiltonian simulation by quantum signal processing," *Phys. Rev. Lett.*, vol. 118, p. 010501, Jan 2017. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.118.010501

[19] ——, "Hamiltonian Simulation by Qubitization," *Quantum*, vol. 3, p. 163, Jul. 2019. [Online]. Available: https://doi.org/10.22331/q-2019-07-12-163

[20] A. Gilyén, Y. Su, G. H. Low, and N. Wiebe, "Quantum singular value transformation and beyond: exponential improvements for quantum matrix arithmetics," in *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, ser. STOC 2019. Association for Computing Machinery, 2019.

[21] S. Bravyi and A. Kitaev, "Universal quantum computation with ideal clifford gates and noisy ancillas," *Phys. Rev. A*, vol. 71, p. 022316, Feb 2005. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevA.71.022316

[22] E. T. Campbell, B. M. Terhal, and C. Vuillot, "Roads towards fault-tolerant universal quantum computation," *Nature*, vol. 549, no. 7671, p. 172–179, Sep. 2017. [Online]. Available: http://dx.doi.org/10.1038/nature23460

[23] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer," *SIAM Journal on Computing*, vol. 26, no. 5, p. 1484–1509, Oct. 1997. [Online]. Available: http://dx.doi.org/10.1137/S0097539795293172

[24] QIR Alliance, *QIR Specification*, 2021, see also https://qir-alliance.org. [Online]. Available: https://github.com/qir-alliance/qir-spec

[25] L. Burgholzer, D. Haag, Y. Stade, D. Rovara, P. Hopf, and R. Wille, "The MQT Compiler Collection: A blueprint for a future-proof quantum-classical compilation framework," in *Design, Automation and Test in Europe*, 2026.

[26] J. contributors. (2025) Jeff Bridges Compilation. [Online]. Available: https://github.com/unitaryfoundation/jeff

[27] C. Lattner, M. Amini, U. Bondhugula *et al.*, "MLIR: Scaling compiler infrastructure for domain specific computation," in *2021 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, 2021, pp. 2–14.

[28] C. Lattner and V. Adve, "LLVM: A compilation framework for lifelong program analysis and transformation," in *CGO*, San Jose, CA, USA, Mar 2004, pp. 75–88.

[29] D. Ittah, T. Häner, V. Kliuchnikov, and T. Hoefler, "QIRO: A static single assignment-based quantum program representation for optimization," *ACM Transactions on Quantum Computing*, vol. 3, no. 3, p. 1–32, Jun. 2022. [Online]. Available: http://dx.doi.org/10.1145/3491247

[30] D. Ittah, J. Fraser, J. Izaac, and O. Di Matteo, "Constant-time hybrid compilation of shor's algorithm with quantum just-in-time compilation," 2025. [Online]. Available: https://arxiv.org/abs/2504.12449

[31] T. Häner, T. Hoefler, and M. Troyer, "Assertion-based optimization of quantum programs," *Proc. ACM Program. Lang.*, vol. 4, no. OOPSLA, Nov. 2020. [Online]. Available: https://doi.org/10.1145/3428201

[32] M. Amy and J. Lunderville, "Linear and non-linear relational analyses for quantum program optimization," *Proc. ACM Program. Lang.*, vol. 9, no. POPL, Jan. 2025. [Online]. Available: https://doi.org/10.1145/3704873

[33] U. Bondhugula, A. Hartono, J. Ramanujam, and P. Sadayappan, "Pluto: A practical and fully automatic polyhedral program optimization system," in *Proceedings of the ACM SIGPLAN 2008 Conference on Programming Language Design and Implementation (PLDI 08), Tucson, AZ (June 2008). Citeseer*, vol. 146, 2008.

[34] S. Verdoolaege, J. Carlos Juega, A. Cohen, J. Ignacio Gomez, C. Tenllado, and F. Catthoor, "Polyhedral parallel code generation for cuda," *ACM Transactions on Architecture and Code Optimization (TACO)*, vol. 9, no. 4, pp. 1–23, 2013.

[35] R. Nigam, S. Thomas, Z. Li, and A. Sampson, "A compiler infrastructure for accelerator generators," in *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, 2021, pp. 804–817.

[36] C. Kim, P. Li, A. Mohan, A. Butt, A. Sampson, and R. Nigam, "Unifying static and dynamic intermediate languages for accelerator generators," *Proceedings of the ACM on Programming Languages*, vol. 8, no. OOPSLA2, pp. 2242–2267, 2024.

[37] T. Neumann, "Efficiently compiling efficient query plans for modern hardware," *Proceedings of the VLDB Endowment*, vol. 4, no. 9, pp. 539–550, 2011.

[38] T. Grosser, J. Ramanujam, L.-N. Pouchet, P. Sadayappan, and S. Pop, "Optimistic delinearization of parametrically sized arrays," in *Proceedings of the 29th ACM on International Conference on Supercomputing*, 2015, pp. 351–360.

[39] R. C. Rocha, P. Petoumenos, B. Franke, P. Bhatotia, and M. O'Boyle, "Loop rolling for code size reduction," in *2022 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*. IEEE, 2022, pp. 217–229.

[40] J. W. Davidson and S. Jinturkar, "Aggressive loop unrolling in a retargetable, optimizing compiler," in *International Conference on Compiler Construction*. Springer, 1996, pp. 59–73.

[41] A. W. Cross, L. S. Bishop, J. A. Smolin, and J. M. Gambetta, "Open quantum assembly language," 2017. [Online]. Available: https://arxiv.org/abs/1707.03429

[42] M. Koch, A. Borgna, S. Sivarajah, A. Lawrence, A. Edgington, D. Wilson, C. Roy, L. Mondada, L. Heidemann, and R. Duncan, "Hugr: A quantum-classical intermediate representation," 2025. [Online]. Available: https://arxiv.org/abs/2510.11420

[43] QuEra Computing Inc, "Kirin: Kernel intermediate representation infrastructure," 2025. [Online]. Available: https://queracomputing.github.io/kirin/latest/

[44] A. Ransford, M. S. Allman, J. Arkinstall, J. P. C. III, and S. F. C. et al., "Helios: A 98-qubit trapped-ion quantum computer," 2025.