Quantum Circuit Equivalence Checking: A Tractable Bridge From Unitary to Hybrid Circuits

Jérome Ricciardi[1,2][0009−0001−7433−8384], Sébastien Bardin[1][0000−0002−6509−3506], Christophe Chareton[1][0000−0001−7113−563X], and Benoît Valiron[2][0000−0002−1008−5605]

[1] Université Paris-Saclay, CEA, List, F-91120, Palaiseau, France, `first.name@cea.fr`
[2] Université Paris-Saclay, CNRS, CentraleSupélec, ENS Paris-Saclay, Inria, Laboratoire Méthodes Formelles, 91190, Gif-sur-Yvette, France, `benoit.valiron@lmf.cnrs.fr`

Abstract. Equivalence checking of hybrid quantum circuits is of primary importance, given that quantum circuit transformations are omnipresent along the quantum compiler chain. While some approaches exist for automating this task, most focus on the simple case of unitary circuits. At the same time, real quantum computing requires hybrid circuits equipped with measurement operators. Moreover, the few approaches targeting the hybrid case are limited to a restricted class of problems. We propose tackling the Quantum Hybrid Circuit Equivalence Checking problem through lifting unitary circuit verification using a transformation known as deferred measurement. We show that this approach alone significantly outperforms prior work, and that, with the addition of specific unitary-level techniques we call separation and projection, it can handle much larger classes of hybrid circuit equivalence problems. We have implemented and evaluated our method over standard circuit transformations such as teleportation, one-way measurement, or the IBM Qiskit compiler, demonstrating its promises. As a side finding, we have identified and reported several unexpected behaviours with the Qiskit compiler.

Keywords: Hybrid classical/quantum circuits, Automated verification, Formal certification.

1 Introduction

Quantum computation is a realm where information is stored on the state of objects governed by the laws of quantum physics [33]. This model of computation is believed to provide important speedup for many applications, ranging from high-performance computing to optimisation. In recent years, quantum computers have become a near-term physical, industrial, and economic reality.

Compared to classical information, quantum data is very peculiar: it cannot be duplicated, and reading is a probabilistic operation done through *measurement*, changing the global state of the memory. Unlike classical data, whose

typical semantics is based on discrete data structures, quantum information is modelled with vectorial structures in Hilbert spaces. Because they manipulate data structures radically different from the classical case, quantum computers require specific developments at every development stage (user languages, verification, optimisation, compilation, etc.) [30,13].

The typical execution flow for quantum programs relies on the notion of *quantum coprocessor*: the quantum memory is stored in an external device, seen as a coprocessor to a CPU, similar to what happens for a GPU, for instance. The quantum coprocessor keeps the memory alive, while the main computation occurs on the classical CPU. A quantum process, therefore, consists of a (classical) interaction with the coprocessor by sending a series of instructions to initialise, update (with *quantum, unitary gates*), and measure the quantum memory to retrieve classical information. A measurement's result is a classical piece of information that might be stored for later use or discarded. Such a series of instructions is represented by a *quantum circuit*. A circuit can generally mix qubit initialisations, unitaries, measurements, and discards. We call *unitary* (or *purely quantum*) a circuit that consists only of quantum gates. A circuit with unitary gates, initialisation, measurements, discards, and classically controlled instruction is called a *hybrid circuit*.

Known as the *deferred measurement principle*, a result from folklore states that measurements and discards can be postponed to the end of the computation. This principle is partly why circuits found in most quantum algorithms are given uniquely in terms of unitary gates: measurement can always be thought of happening at the end of the computation. However, from an implementation point of view, it makes sense to consider the hybrid case as many quantum processes, such as repeat-until-success, measurement-based quantum computation, or optimisation techniques, rely on it.

Quantum circuit transformations turn a circuit into another (equivalent) quantum circuit. Such techniques are key components of the quantum software stack, whether for optimisation, adaptation to hardware capabilities in terms of qubits and connectivity, error correction, circuit robustification, one-way measurement. Compilers like *Qiskit* do propose several transformation passes. As these passes are omnipresent between the programmer and the quantum hardware, checking their correctness is paramount. While we could envision quantum compilers fully certified in interactive theorem provers, another approach involves designing dedicated automated circuit equivalence checkers [30].

This paper focuses on the automated equivalence verification problem arising from quantum circuit transformation: how to verify that two quantum circuits are functionally equivalent. In particular, we focus on verifying circuits involving initialisation, measurements, discards, and classically controlled instructions: hybrid circuits.

Challenge with hybrid circuits equivalence verification. Current methodologies for automatic equivalence checking predominantly focus on purely quantum, unitary circuits [1,2,29,36]. This restriction is becoming increasingly unrealistic as physical chips progress. Indeed, the current trend in the design of quantum

compilation toolchains tends to split the purely quantum processes into smaller components, to make them more robust to noise. Moreover, even pure quantum primitives require hybridisation and classical control for error detection and correction.

Compared to purely quantum circuits, the formalisation of hybrid computations brings several extra difficulties.

A first difficulty is that hybrid quantum computation is, by nature, nondeterministic: the functional behaviour of a quantum program is a branching probabilistic structure instead of a linear trace. Quantum computation can then be regarded as a strict superset of probabilistic computation—a field where the formal analysis of programs is still a research question [6].

A second difficulty is that a quantum program manipulates both quantum and classical registers. Some of these registers might contain "garbage", i.e. data that should not be considered as output: a relevant equivalence notion should not consider them. We talk of *partial equivalence* when the state equivalence is evaluated only up to non-discarded registers, an essential feature to model a program's state and its observable behaviour. However, discarding quantum information is not innocuous, as it corresponds to tracing it out: discarding corresponds to a measurement and is then a probabilistic process. Therefore, the question of *partial equivalence* of quantum programs is not trivial to define.

Finally, the measurement operation induces a deep conceptual shift from a (quantum) deterministic process purely modelled in linear algebra to a stochastic process acting on a mix of quantum states and usual data structures (Boolean values, integers, etc).

Designing a unified, tractable mechanism for the equivalence of hybrid quantum circuits is therefore a major challenge for formal methods.

Contributions. To address this challenge, this paper presents an approach based on two main ingredients: (1) a classification of hybrid circuit equivalence problem instances, and (2) the lifting of verification methods designed for the unitary case to verify hybrid circuits, based on the deferred measurement principle mentioned earlier. In doing so, we advocate for the versatility of the formalism of *path-sums* in the context of hybrid quantum circuit equivalence checking. More precisely, we bring the following contributions:

1. We clarify the current landscape of hybrid circuit equivalence checking (Section 3). In particular, we draw the separation between circuits with and without discards and the induced partitioning of equivalence cases.
2. We propose a generic method based on deferred measurement that extends unitary circuit equivalence checkers to support *hybrid circuits* (Section 3).
3. We introduce a verification tool, SQbricks, dedicated to hybrid circuit equivalence checking (Section 4), included when discard is at stake. We systematically evaluate this implementation.

Overall, while our technique already allows for significant clarification and pushes forward the state-of-the-art of hybrid circuit equivalence checking, we also believe

that our findings and benchmark establish a solid baseline for future research in the field. To this end, *SQbricks and our experimental setup will be made available as open source.*

State of the Art. Equivalence checking of purely unitary circuits is prevalent among the state-of-the-art. One can cite AutoQ [17,16,1], Feynman [2,3,4], SliQEC [41,15], and PyZX [29]. These tools leverage distinct methodologies, respectively: Tree Automata, Path-Sum, Decision Diagrams, and ZX-calculus.

On the other hand, to the best of our knowledge, only QCEC [10] and VeriQC [24] address the equivalence checking of hybrid quantum circuits. However, concrete implementations are limited to the case without discard (characterized as the DisFree case in Section 4.5 of [24]) .

QCEC relies on a combination of ZX-calculus, Decision Diagrams, Simulation, and deferred measurement for hybrid circuit equivalence. Yet, the approach is limited [10, Section. 4, p. 4] to the narrow case of hybrid circuits with no discards, preventing the analysis of typical transformations such as teleportation, one-way measurement, error correction, optimisations with ancillas, etc., and any form of non-reversible computation.

VeriQC [24] implements measurement and classical control in quantum circuits using Tensor Decision Diagrams. The authors have already identified one of the main challenges in hybrid circuit equivalence verification: managing discards, and they have formally defined an equivalence relation that accounts for it. Still, their tool does not tackle it and does not adress discard.

Table 1 provides an overview of the current state-of-the-art and our results. It details the tools in terms of the technology they employ, the types of equivalence verification they support, and whether they rely on projection and separation techniques (last column).

2   Background

This section is devoted to the presentation of the model of quantum computation, in particular, its main programming model: quantum circuits. We then introduce the two main aspects of the state-of-the-art that underpin our proposal: the *deferred measurement principle* and the *path-sums* semantics.

Quantum Circuits. The interaction with the quantum coprocessor consists of emitting a sequence of initialisation, unitary gates, measurements, and discards, potentially classically controlled by the result of previous measurements.

Example 1. As an example of a hybrid circuit, Figure 1a illustrates the *teleportation protocol* [7]. This protocol operates over mixed quantum and classical data and describes the transmission of a quantum state from Alice (input register InpVar) to Bob (observable register Obs). The protocol highlights the general structure of quantum circuits: initialisation (Init), unitary evolution, measurement, discards, and classically controlled operations. The circuit consists of three quantum wires (single lines) and two classical wires (double lines). It employs

Table 1: Automatic Quantum Circuit Equivalence Tools. SQV: SQbricks-Verif, SQL: SQbricks-Lifting, (*): Any unitary equivalence checker considered here, **: The approach is unclear on these points (VeriQC), †: Meet our prerequisites (featuring separation tests and projections) to be able to verify lifted problems from Mix and Dis, ✓: Valid, ×: Invalid.

Name	Techno.	Equivalence Checking Abilities				Prerequisites†
		Unitary	Hybrid			
			DisFree	Mix	Dis	
AutoQ [17,1]	Tree Automata	✓	×	×	×	×
AutoQ-2 [16]	Tree Automata	✓	×	×	×	×
Feynman [2,3,4]	Path-Sum	✓	×	×	×	×
PyZX [29]	ZX Calculus	✓	×	×	×	×
SliQEC [41,15]	Decision Diagram (DD)	✓	×	×	×	×
QCEC [36,34,10]	ZX, DD, Simulation	✓	✓	×	×	×
VeriQC [25,24]	Tensor DD	✓	✓	**	**	×
SQV	Path-Sum	✓	×	×	×	✓
SQV + SQL	Path-Sum	✓	✓	✓	✓	✓
(*) + SQL		✓	✓	×	×	×

(a) Standard presentation


(b) With deferred measurement

Fig. 1: The one-qubit teleportation algorithm

three types of gates: the Hadamard gate (H), the (controlled) NOT gate (X) and the phase gate ($R_z(1)$). Measurement $\boxed{\nearrow}$ is a non-deterministic, described by the so-called *Born's rule*, with results stored in classical wires $\oplus$. Classical wires are assumed to be initialised at 0. Here, their values are not outputs of the circuit: they are discarded with ══║ . Historically, mathematical semantics for quantum circuits have relied on the underlying mathematical representation of unitary gates: linear, unitary maps. One problem is that the matrix size corresponds to the number of basis elements in the corresponding vector space: it is exponentially costly to compute. Another issue is that quantum circuits generally contain measurements, typically handled with more involved representations such as density matrices and superoperators. By sake of space, we cannot further introduce these quantum circuit component and their interpretation. We refer the desirous reader to [33] (Ch.4) for the standard introductory material.

The deferred measurement is a circuit transformation that postpones all measurements to the end by replacing classical controls with quantum controls, leaving the inner circuit purely unitary. This transformation preserves the semantics, resulting in equivalent circuits. Figure 1b shows the result of applying the process to the teleportation algorithm of Figure 1a. The output circuit consists of a round of initialisations ⊢── , a unitary block, and a round of measurements, possibly immediately followed by discards (represented with ──⊪). We refer to the resulting pattern as the IUM circuits. Note that in our figures, the alignment of the wires in the drawing is not meaningful (they might, for instance, be shuffled inside the circuit). A formal definition of our implementation of the deferred measurement transformation is provided in Appendix B.1.

Path-sums. An alternative to matrix representation has recently been proposed: *path-sums* [2,3,40,39,21,12]. This symbolic representation gives a more compact representation for purely quantum circuits, and has been at the core of an equivalence checking tool for unitary circuits: Feynman [2,3]. The name refers to *Feynman's paths*, where the quantum evolution of a system is represented as a weighted sum of possible "paths". In the path-sums formalism, these weights are parameterised by the input basis states.

The path-sums formalism supports symbolic sequential ([2, Definition 2.6]) and parallel composition of operators. Nevertheless, path-sums are over-expressive: two circuits corresponding to the same linear operation might have distinct path-sum representations. To reconcile such equivalence representations, path-sums come with an equational theory, made of (in addition to Boolean and dyadic ring theories) a set of rewriting rules simplifying path-sum while preserving $\mathcal{V}$ equivalence.

3   Partial Equivalences of Quantum Circuits

The core of this paper focuses on the equivalence checking of quantum circuits. This section is dedicated to presenting our proposition's context and main structuring elements. For the sake of readability, in the following we assume the particular case where all input wires are quantum (the general case, with possibly classical inputs, is formally treated in Appendix B). The prevalent approach to equivalence checking within the field predominantly addresses circuits with no discards. This paper elaborates on our methodological approach to adapt and extend equivalence checking to encompass circuits with discards.

Shapes of Hybrid Circuits. We claim the presence or the absence of discards to be the main difficulty in the equivalence checking of hybrid circuits. It straightforwardly generates a three classes typology of hybrid circuit equivalence instances:

– DisFree concerns circuits without discards. The primary applications of DisFree involve dynamising unitary circuits: basically, the transformation inverse of the deferred measurement, introducing measurement instructions in order to turn quantumly controlled commands into classically controlled ones. A

key usage is robustification, with examples including the Quantum Fourier Transform [11] and Quantum Phase Estimation [18]. The state-of-the-art methods for hybrid circuits currently address confines to this case [10,24].

- Dis, where both circuits feature discard. Dis is the most generic equivalence task over hybrid circuits teleportation, one-way measurement, error correction, optimisations with ancillas, etc. In fact, most of the non-reversible quantum processes require the use of ancillas and discards.
- Mix is the mixed case, where one circuit is with discards, and the second is without. With Mix, we can verify the equivalence between unitary and hybrid circuits resulting from transformations such as one-way measurement or teleportation.

Taking advantage of the Deferred Measurement Transformation. To open the field to the equivalence checking of classes Mix and Dis, we propose a unified approach by capitalising on the *deferred measurement principle* discussed in Section 2.

Given a circuit C, the deferred measurement transformation isolates the initialisation, unitary, and measurement components, obtaining a IUM circuit. A central result of this paper consists in characterising the relationship between the equivalence of two circuits C_1 and C_2 possibly with measures and discards, and the equivalence of their deferred measurement versions $I_1 U_1 M_1$ and $I_2 U_2 M_2$. Indeed, it can be shown that if (i) I_1 (resp. M_1) and I_2 (resp. M_2) are equal and (ii) the unitary blocks U_1 and U_2 behaves equivalently over non-discarded qubits and after initialisation (unitary circuit partial equivalence), one can directly infer the hybrid equivalence between $I_1 U_1 M_1$ and $I_2 U_2 M_2$, and thus address the original problem.

Hence, while our method pre-processes hybrid circuits along the deferred measurement transformation, it does not properly apply the principle. Instead, it builds a logical over-approximation of the equivalence relation that takes advantage of the deferred measurement transformation (the IUM circuit normal form) while eventually ignoring the final measurement.

4   Implementation

We have implemented our method in a prototype named SQbricks, made of about 3,500 lines of OCaml code. The code will be made open source[3]. The implementation comprises two parts: (1) the SQbricks-Verif component implements a path-sum calculus to perform unitary verification and partial equivalence checking; (2) the SQbricks-Lifting component focuses on our generic lifting method, based on the deferred measurement transformation.

SQbricks-Verif (SQV) is a unitary circuit verification tool based on path-sums, with a core rewriting system. Circuits are built from a pseudo-universal set of

---

[3] https://github.com/Qbricks/qbricks.github.io/tree/main/Artifacts/SQbricks

gates, comprising H, X, U1 (Z rotation up to global phase) and combined through sequence and quantum control. Furthermore, the user can explicitly indicate the address of qubits that are discarded.

SQbricks-Lifting (SQL) implements deferred measurement to extend unitary equivalence verification tools to hybrid circuits. This approach maintains consistency between hybrid circuit classes (DisFree, Mix, and Dis) while ensuring broad interoperability with existing verification tools through the OpenQASM 2 standard.

5   Experimental evaluation

To assess our circuit equivalence checking technique, we conducted a series of benchmarks, all led on a laptop using Linux Mint 21.2, equipped with an Intel Core i9-9880H CPU (2.3GHz x 8), 31 GiB of RAM, and a 500 GB SSD.

5.1   Experimental setup

Considered tools. We evaluate SQbricks-Lifting with the state-of-the-art unitary circuit verification tools AutoQ (Tree Automata), Feynman [2,22] (Path-Sums), PyZX [29,35] (ZX) and QCEC [8,9,32] in unitary mode (ZX, Decision Diagram, Simulation, labelled *lifting*, L in the following tables) in addition to SQbricks-Verif (Path-Sums). We compare our overall performance for hybrid equivalence checking against QCEC [10] (in full hybrid mode, labelled *standalone*, S)and VeriQC [24] (figures from Table 1 in [24], since we were not able to use the tool directly (Appendix, Section A.2).

Circuit collection. We selected two libraries of quantum circuits in OpenQASM 2 [19] commonly used in the literature—VeriQBench [14] and the Feynman library [2], plus a hybrid implementation of the Shor algorithm [37] (Hybrid Circuit) inspired by the version provided in QASMBench [31]. For each library, we only retained the circuits that SQbricks could parse (some seem to have issues or need refinement, see Appendix A.1). Observe that multiple samples represent the same algorithms with varying parameter settings, typically the size. This results in a compilation of 420 unitary circuits and 204 hybrid circuits.

Circuit equivalence challenges. We consider the following circuit transformations. (1) **Qiskit$_{tr}$**: the `generate_preset_path_manager` of IBM Qiskit [28] with an optimisation level of 3 (the highest level) to maximise circuit transformations. (2) **OWM**: One-Way Measurement [20], a circuit transformation aiming at enhancing the practical application of quantum computers by minimising the time a qubit remains coherent. (3) **Tele**: Teleportation, described in Figure 1a. In the end, our equivalence challenge categories are:

- DisFree, with two classes of challenges: 169 pairs of unitary-hybrid circuits provided by VeriQBench and implementing the same algorithms (U-DC), either as a unitary or as a hybrid circuit; and 88 pairs consisting of a hybrid circuit and its transformed version using $Qiskit_{tr}$ (DC-Q(DC)).

- Mix, with three classes of challenges: 347 pairs made of a unitary circuit together with its OWM version (O(U)-U); 406 pairs made of a unitary circuit and its Tele version (T(U)-U); 88 pairs corresponding to a hybrid circuit transformed by both $Qiskit_{tr}$ and OWM (O(DC)-Q(DC)).
- Dis: 347 pairs (O(U)-T(U)) obtained from unitary circuits transformed with OWM and Tele.

5.2   Experimental observations

We address the three following research questions.

- RQ1: Can we lift unitary verification for checking DisFree equivalence?
- RQ2: How does it compare to prior works on DisFree equivalence?
- RQ3: Can we lift unitary verification for checking Dis and Mix equivalence?

We also performed a sanity check to evaluate the correctness of the considered tools (Section 5.3), and report some additional findings about unexpected behaviours in $Qiskit$ we found along our experiments (Section 5.3).

RQ1: Can we lift unitary verification for checking DisFree?

For DisFree challenges, we could draw performance comparisons between SQV and our selected set of unitary verification tools. To do so, we performed a pre-processing deferred measurement transformation over DisFree circuits. Results are summarized in Table 2: path-sum based methods (Feynman and SQV) achieved perfect success rates. In contrast, PyZX demonstrated limitations, failing to verify equivalence between two CSWAP variants (Appendix, Example 2) and showing scalability issues. For instance, PyZX took 534s to verify a DC-Q(DC) QPE instance of size 35, whereas Feynman verified an instance of size 42 in just 3s. AutoQ successfully verified all circuit families for small instances but faced scalability challenges, such as being limited to DC-Q(DC) to QPE of size 6.

Conclusion: Our approach successfully lifts unitary verification tools to handle hybrid DisFree cases, with path-sums showing superior performance compared to ZX-calculus and automata-based methods.

RQ2: How does it compare to prior works on DisFree? We compare QCEC [10] and VeriQC [24] for DisFree hybrid equivalence, evaluating QCEC directly and VeriQC via published results. QCEC (standalone), lifted Feynman, and lifted SQV show similar performance on 88 DC-Q(DC) challenges, but QCEC fails all 169 U-DC challenges, and incorrectly reports 27 non-equivalence proofs, indicating correctness issues (see Appendix, Example 3 for a minimal example). VeriQC [24] performs hybrid equivalence checking on DisFree. [4] Unfortunately, despite contacting VeriQC [24] authors, we couldn't use this tool for our experiments. We compared our methods against its published [24] experimental performance

---

[4] The paper claims results over Dis [24, Definition 2]. However, [24, Section 4.5] explains how to check dynamic circuit equivalence with TDDs, but is limited to circuits without discard (our DisFree).

Table 2: Evaluation results of the lifting application of deferred measurement. S: Standalone, L: Lifted, DC: Dynamic Circuit, subclass of Hybrid Circuit without Discard, O: OWM,   T: Tele,   Q: $Qiskit_{tr}$, TO: Time Out (10 min) Wrong: equivalence check returns not equivalent for equivalent circuits, NA: Not Applicable, NW: Not Working.

| | | Success | | | | | | | |
| | | DisFree | | Mix | | | Dis | | |
Tool	Lift	U-DC	DC-Q(DC)	O(U)-U	T(U)-U	O(DC)-Q(DC)	O(U)-T(U)	Total	TO
QCEC	S	0 Wrong: 27	87	NA	NA	NA	NA	87 Wrong: 27	143
VeriQC (Cf. Table 3)	S	NW	NW	NA	NA	NA	NA	NA	NA
AutoQ-2.0	L	10	15	NA	NA	NA	NA	25	23
Feyn-24	L	169	88	NA	NA	NA	NA	257	0
PyZX	L	73	66	NA	NA	NA	NA	139	112
QCEC	L	151	69	NA	NA	NA	NA	220	19
SQV	L	169	88	180	343	50	137	967	426
#challenges		169	88	347	406	88	347	1445	

Table 3: Comparing against VeriQC from published results on their available benchmark (selection of the most significant results, time in seconds, ✓: Success)

| | | SQbricks-Lifting + | | | | |
Task	VeriQC [24]	AutoQ	Feyn-24	PyZX	QCEC	SQV
QFT _11	0.86 ✓	Err	0,01 ✓	0,39 ✓	0,01 ✓	0,01 ✓
QFT _16	23.38 ✓	Err	0,02 ✓	3,64 ✓	0,02 ✓	0,03 ✓
pe_9	1.62 ✓	Err	0,01 ✓	0,29 ✓	0,01 ✓	0,01 ✓
phaseflip	0.18 ✓	2,50 ✓	0,00 ✓	0,01 ✓	0,00 ✓	0,00 ✓
teleportation	0.01 ✓	0,00 ✓	0,00 ✓	0,00 ✓	0,00 ✓	0,00 ✓
state_inj_T	0.01 ✓	0,00 ✓	0,00 ✓	0,00 ✓	0,00 ✓	0,00 ✓

data. Results are shown in Table 3. With our lifting approach, all tools except AutoQ verify the entire VeriQC benchmark faster than VeriQC itself.

Conclusion: Our lifting method is highly effective on DisFree challenges. While our method matches QCEC's performance on its best subcategory, it outperforms both QCEC and VeriQC in all other cases.

RQ3: Can we lift unitary verification for checking Dis and Mix?

This problem considers challenges from Dis and Mix, out of the scope of prior work. Therefore, no comparative analysis was possible here. Our method effectively addresses a substantial portion of the Mix and Dis hybrid equivalence challenges, handling 39.5% of Dis and 67.7% of Mix.

Conclusion: Our approach can indeed address hybrid equivalence checking out of the scope of the current state-of-the-art tools, with a reasonable success rate, establishing an acceptable first solution for these problems.

5.3   Additional Findings

Sanity check. We also performed a sanity check consisting of 73 equivalence tasks from VeriQBench and QASMBench, with deliberately modified quantum circuits (mutants). These modifications ensure non-equivalence by design. All versions of Feynman and SQbricks successfully passed the sanity check with no false positives. Other tools exhibited two types of failures, primarily related to rotation gates: (1) PyZX and QCEC failed to handle very small angles ($\leq \pi/2^{26}$ and $\leq \pi/2^{27}$ respectively). (2) AutoQ failed to distinguish between controlled and uncontrolled Z axis rotation gates ($\mathsf{CR_z(k)}$ and $\mathsf{R_z(k)}$ gates).

Unexpected behaviors with $Qiskit_{tr}$. During our experiments, we uncovered two bugs in the Qiskit compiler: (1) **Angle approximation:** Qiskit version 1.1.0 approximated small angles to 0, leading to incorrect circuit simplifications in e.g., QFT, QPE, Shor algorithm, etc, when quantum registers are over 42 qbits[5]. This issue was fixed in version 1.4.0. (2) **Introduction of Floats:** In version 1.4.0, transformations introduced floating-point values (e.g., converting $3\pi/32$ to 0.2945243112740431) instead of rational numbers, causing equivalence loss in some circuits. We reported this to the Qiskit team. These findings highlight the practical utility of our method in identifying and mitigating approximation-related issues in quantum circuit transformations.

References

1. Abdulla, P.A. *et al.*: Verifying quantum circuits with level-synchronized tree automata. Proc. POPL **9** (2025).
2. Amy, M.: Towards Large-scale Functional Verification of Universal Quantum Circuits. EPTCS **287** (2019).
3. Amy, M.: Complete equational theories for the sum-over-paths with unbalanced amplitudes. EPTCS **384** (2023).
4. Amy, M., Lunderville, J.: Linear and non-linear relational analyses for quantum program optimization. Proc. POPL **9** (2025).
5. AutoQ Git, https://github.com/fmlab-iis/AutoQ
6. Barthe, G., Katoen, J.P., Silva, A. (eds.): Foundations of Probabilistic Programming. Cambridge University Press (2020)
7. Bennett, C.H. et al.: Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. Phys. Rev. Lett. **70**, 1895–1899 (1993).
8. Burgholzer, L., Wille, R.: Improved DD-based equivalence checking of quantum circuits. In: Proc. ASP-DAC (2020).

---

[5] This simplification imposed the induced restrictions over our experiments

9. Burgholzer, L., Wille, R.: The power of simulation for equivalence checking in quantum computing. In: Proc. DAC (2020).
10. Burgholzer, L., Wille, R.: Handling Non-Unitaries in Quantum Circuit Equivalence Checking. In: Proc. DAC (2022).
11. Bäumer, E. *et al.*: Quantum Fourier transform using dynamic circuits (2024).
12. Chareton, C. *et al.*: An automated deductive verification framework for circuit-building quantum programs. ESOP (2021).
13. Chareton, C. *et al.*: Formal methods for quantum algorithms. In: Handbook of Formal Analysis and Verification in Cryptography. CRC Press (2023).
14. Chen, K. *et al.*: Veriqbench: A benchmark for multiple types of quantum circuits. arXiv:2206.10880 (2022).
15. Chen, T.F. *et al.*: Partial equivalence checking of quantum circuits. QEC 2022.
16. Chen, Y.F. *et al.*: AutoQ 2.0: From verification of quantum circuits to verification of quantum programs. arXiv:2411.09121 (2024).
17. Chen, Y.F. *et al.*: AutoQ: An automata-based quantum circuit verifier. CAD 2023.
18. Córcoles, A.D. *et al.*: Exploiting dynamic quantum circuits in a quantum algorithm with superconducting qubits. Phys. Rev. Lett. **127**, 100501 (2021).
19. Cross, A.W., Bishop, L.S., Smolin, J.A., Gambetta, J.M.: Open quantum assembly language. arXiv:1707.03429 (2017).
20. Danos, V., Kashefi, E., Panangaden, P.: The Measurement Calculus. (2007)
21. Deng, H., Tao, R., Peng, Y., Wu, X.: A case for synthesis of recursive quantum unitary programs. Proc. POPL (2024).
22. Feynman Git, `https://github.com/meamy/feynman`
23. Devitt, S. J. and Munro, W. J. and Nemoto, K.: Quantum error correction for beginners. (2013),
24. Hong, X., Feng, Y., Li, S., Ying, M.: Equivalence checking of dynamic quantum circuits. In: Proc. ICCAD (2022)
25. Hong, X. *et al.*: A tensor network based decision diagram for representation of quantum circuits. ACM Trans. Des. Autom. Electron. Syst. **27**(6) (Jun 2022).
26. Ying M. and Zhou L. and Barthe G.: Laws of Quantum Programming. (2024).
27. Ioannou, L. M.: Computational complexity of the quantum separability problem (2006).
28. Javadi, A. *et al.*: Quantum computing with Qiskit. arXiv:2405.08810 (2024).
29. Kissinger, A., van de Wetering, J.: PyZX: Large Scale Automated Diagrammatic Reasoning. In: Proc. QPL (2020).
30. Lewis, M., Soudjani, S., Zuliani, P.: Formal verification of quantum programs: Theory, tools, and challenges. ACM TQC **5**(1) (2023).
31. Li, A., Stein, S., Krishnamoorthy, S., Ang, J.: Qasmbench: A low-level qasm benchmark suite for NISQ evaluation and simulation. arXiv:2005.13018 (2022).
32. MQT-QCEC, `https://mqt.readthedocs.io/projects/qcec/en/latest/`
33. Nielsen, M.A., Chuang, I., Grover, L.K.: Quantum computation and quantum information. Am. J. Ph. **70**(5), 558–559 (2002).
34. Peham, T., Burgholzer, L., Wille, R.: Equivalence checking of quantum circuits with the ZX-calculus. IEEE JESTCS **12**(3), 662–675 (2022).
35. PyZX. `https://pyzx.readthedocs.io/en/latest/gettingstarted.html`
36. Sander, A., Burgholzer, L., Wille, R.: Equivalence checking of quantum circuits via intermediary matrix product operator. arXiv:2410.10946 (2024).
37. Shor, P.W.: Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. SIAM J. Comp. **26**(5), 1484–1509 (1997).
38. VeriQC dynamic quantum circuits benchmarks, `https://github.com/Veriqc/EC-for-Dynamic-Quantum-Circuits/tree/main/Benchmarks2`

39. Vilmart, R.: The structure of sum-over-paths, its consequences, and completeness for Clifford. arXiv:2003.05678 (2020).
40. Vilmart, R.: Rewriting and completeness of sum-over-paths in dyadic fragments of quantum computing. LMCS **20(1)** (2024).
41. Wei, C.Y. *et al.*: Accurate BDD-based unitary operator manipulation for scalable and robust quantum circuit verification. In: DAC (2022).

A   Experimental evaluation details

This appendix provides details about our experimental evaluation, including specifics about the collection of circuits we used (Appendix A.1) and information about the tools we compared ourselves to (Appendix A.2).

A.1   Details on Circuit Collection

Our circuit collection comprises two categories:

- **Unitary Circuits**: 420 circuits, including $43/44$ circuits from the Feynman library [6]. and $377/782$ circuits from the VeriQBench library combinational subset, and without the sequential and variational subsets due to parsing issues. Parsing issues arise if a circuit is ill-formed or contains elements not accounted for in our syntax, such as macros.
- **Hybrid Circuits**: 204 circuits, including $198/205$ circuits from the VeriQBench library (QPE and QFT, bit flip and phase flip correction, state injection and teleportation) and an implementation of Shor [37] over 5 qubits inspired from the QASMBench [31] library.

A.2   Tool Limitations and Failure Cases

This appendix provides information on the benchmark for VeriQC (Table 3), and presents concrete examples demonstrating specific limitations of state-of-the-art quantum circuit verification tools, including cases where tools return incorrect or inconclusive results despite the functional equivalence of the circuits being compared.

VeriQC. After encountering difficulties in reproducing their experiments with DisFree, we began email correspondence with the authors. However, we faced on-going challenges in replicating their results. As a result, we utilised their findings: [24, Table 1].

Example 2 (PyZX minimal inconclusive result). The following pair of circuits implements the controlled-swap operation (CSWAP) in two different ways. This operation swaps two qubits conditionally on the state of a third control qubit. For basis states $|x_0, x_1, x_2\rangle$, the operation should:

---

[6] Except cycle_17_3.qasm which has an implementation issue

- Leave x_0 unchanged (control qubit)
- Swap x_1 and x_2 when $x_0 = 1$
- Leave x_1 and x_2 unchanged when $x_0 = 0$

While these implementations are functionally equivalent (both implement controlled-swap), PyZX [29] (version 0.9.0) fails to verify this equivalence, demonstrating a limitation in handling certain control structures.

Fig. 2: CSWAP first implementation   Fig. 3: CSWAP alternative implementation

Example 3 (QCEC minimal incorrect result). The circuits below illustrate a minimal case where QCEC [10] (version 2.8.1), when using its deferred measurement option, incorrectly determines that the circuits C_1 and C_2 are not equivalent.

Fig. 4: C_1 Fig. 5: C_2

B   Technical Details

B.1   Lifting unitary verification tools for hybrid circuits

This section provides the main formal ingredients for our deferred measurement-based lifting of unitary circuits' equivalence to the hybrid case.

Definition 1 (Hybrid circuits). *A hybrid circuit is a sequence of instructions generated by the following syntax*

$$G := Ph(k) \mid R_z(k) \mid X \mid H$$
$$GA := Apply(G, [qb], [qb])$$
$$Ins := GA \mid if\ cb\ then\ GA \mid \mathsf{Meas}(qb, cb) \mid \mathsf{Init}(qb) \mid Not(cb)$$
$$C := [Ins]$$

where $[t]$ denotes a standard list construction of type t, built either as the nill list $[l] := $ nill or as a t type object o appended to another t list $[l] := o - [l']$. By abuse of notation, we use the same $l - l'$ for l, l' being either a list or a type t object (assimilated to a list of one single element).

Apply$(G, [qb_1], [qb_2])$ intuitively commands the parallel application of gate G on qubits in $[qb_2]$, controlled by the conjunction of qubits in $[qb_1]$; $\mathsf{Meas}(qb, cb)$ commands the measure of qubit qb and the storage of the obtained data in classical bit cbA circuit is well-formed *if it respects the following syntactical constraints: that (i) a classical wire should receive at most one measurement result in a given circuit and (ii) a quantum wire is not further addressed when having been measured.*

The deferred measurement principle is introduced in [33] as a *rather obvious* property of circuits: that measurement can always be moved from an intermediate stage of the computation to the end of the circuit, while classically controlled instructions are replaced by quantum conditionals. Surprisingly, to the best of our knowledge, the first formal and generic proof of the principle was established as late as 2022 [23], and a computer-assisted proof was even recently given in [26].

For a circuit C, we write $\texttt{lInit}(C)$ the sum of, for each initialisation command Init, the number of non-initialisation commands that precede it in C. Symmetrically, $\texttt{eMeas}(C)$ denotes the sum of, for each measure command Meas in C, the number of non-measure commands succeeding it in C. We also informally introduce the function m_C. It maps a classical wire cb to a quantum wire qb in instructions. In our deferred measurement circuit transformation, as the measurement Meas(qb, cb) of a qubit is postponed to the end of the execution, the intermediary classical control commands over cb should be turned into quantum control commands over qb. This is achieved by applying the instruction transformation m_C(qb, cb) to the instructions Meas(qb, cb) is permuted with. In addition, it turns a classical bit-flip instruction Not(cb) into its quantum counterpart Apply(X, nill, qb). Formally, for any instruction Ins, we have that:

$$\textsf{Ins}[m_C(\textsf{cb}, \textsf{qb})] = \textbf{if } \textsf{Ins} := \text{if cb then } \textsf{Apply}(\textsf{G}, [co], [ta]) \textbf{ then } \textsf{Apply}(\textsf{G}, [co \cup \{\textsf{qb}\}], [ta])$$
$$\textbf{else if } \textsf{Ins} := \textsf{Not}(\textsf{cb})$$
$$\textbf{then } \textsf{Apply}(\textsf{X}, \textsf{nill}, \textsf{qb})$$
$$\textbf{else } \textsf{Ins}$$

Definition 2 (Deferred measurement circuit transformation). *The deferred measurement transformation is built as a double inductive rewriting pass, where initialisation (resp. measurement) commands systematically commute with non-init (non-measure) commands whenever they occur in non-initial (resp. non-final) position.*

$$
\begin{aligned}
\textit{pI(C)} := \textit{if} \quad & \textit{lInit}(C) = 0 & & \textit{then } C \\
\textit{else if } & C := \textsf{Init}(qb) - C' & & \textit{then } \textsf{Init}(qb) - \textit{pI}(C') \\
\textit{else if } & C := \textit{Ins} - \textsf{Init}(qb) - C' & & \textit{then } \textsf{Init}(qb) - \textit{pI}(\textit{Ins} - C') \\
\textit{else if } & C := \textit{Ins} - C' & & \textit{then } \textit{pI}(\textit{Ins} - \textit{pI}(C'))
\end{aligned}
$$

$$
\begin{aligned}
\textit{lM(C)} := \textit{if} \quad & \textit{lM}(C) = 0 & & \textit{then } C \\
\textit{else if } & C := C' - \textsf{Meas}(qb, cb) & & \textit{then } \textit{lM}(C') - \textsf{Meas}(qb) \\
\textit{else if } & C := C' - \textsf{Meas}(qb, cb) - \textit{Ins} & & \textit{then } \textit{lM}(C' - \textit{Ins}[m_C(qb, qb)]) - \textsf{Meas}(qb, cb) \\
\textit{else if } & C := C' - \textit{Ins} & & \textit{then } \textit{lM}(\textit{lM}(C') - \textit{Ins})
\end{aligned}
$$

$$\textsf{DM}(C) := \quad \textit{lM}(\textit{pI}(C))$$

This transformation provides a circuit equivalent to C and sequentially structured as three successive blocks $[\textsf{Init}]_C - [\textsf{U}]_C - [\textsf{Meas}]_C$ of (i) initialisation (ii) unitary gate application, and (iii) measure commands. Formally, we have the following theorem:

Theorem 1 (Deferred measurement). *Let C be a hybrid circuit, then: (i)* $\textsf{DM}(C) \equiv_H C$, *(ii)* $\textit{lInit}(\textsf{DM}(C)) = 0$, *(iii)* $\textit{eMeas}(\textsf{DM}(C)) = 0$

Proof (Sketch). By structural induction over `C`. Transformation `pI` preserves the semantical equivalence and ensures condition 2., transformation `1M` preserves both and ensures condition 3 in addition.

Example 4. As an illustration of Definition 2, Figure 1b draws it application to the teleportation case discussed in Example 1: measurement instructions are delayed to the end of the execution, and intermediary classical instruction – originally controlling over measurement results – are turned into quantum controlled instruction – controlling over the corresponding *not yet measured* quantum wires.