

Quantum Gate Pattern Recognition and Circuit Optimization for Scientific Applications

Wonho Jang¹, Koji Terashi², Masahiko Saito², Christian W. Bauer³, Benjamin Nachman³, Yutaro Iiyama², Tomoe Kishimoto², Ryunosuke Okubo¹, Ryu Sawada², and Junichi Tanaka²

¹Department of Physics, The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan

²International Center for Elementary Particle Physics (ICEPP), The University of Tokyo, 7-3-1 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan

³Physics Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

There is no unique way to encode a quantum algorithm into a quantum circuit. With limited qubit counts, connectivities, and coherence times, circuit optimization is essential to make the best use of near-term quantum devices. We introduce two separate ideas for circuit optimization and combine them in a multi-tiered quantum circuit optimization protocol called AQCEL. The first ingredient is a technique to recognize repeated patterns of quantum gates, opening up the possibility of future hardware co-optimization. The second ingredient is an approach to reduce circuit complexity by identifying zero- or low-amplitude computational basis states and redundant gates. As a demonstration, AQCEL is deployed on an iterative and efficient quantum algorithm designed to model final state radiation in high energy physics. For this algorithm, our optimization scheme brings a significant reduction in the gate count without losing any accuracy compared to the original circuit. Additionally, we have investigated whether this can be demonstrated on a quantum computer using polynomial resources. Our technique is generic and can be useful for a wide variety of quantum algorithms.

1 Introduction

Recent technology advances have resulted in a variety of universal quantum computers that are being used to implement quantum algorithms.

Koji Terashi: koji.terashi@cern.ch

However, these noisy-intermediate-scale quantum (NISQ) devices [1] may not have sufficient qubit counts or qubit connectivity and may not have the capability to stay coherent for entirety of the operations in a particular algorithm implementation. Despite these challenges, a variety of applications have emerged across science and industry. For example, there are many promising studies in experimental and theoretical high energy physics (HEP) for exploiting quantum computers. These studies include event classification [2–6], reconstructions of charged particle trajectories [7–10] and physics objects [11, 12], unfolding measured distributions [13] as well as simulation of multi-particle emission processes [14, 15]. A common feature of all of these algorithms is that only simplified versions can be run on existing hardware due to the limitations mentioned above.

There are generically two strategies for improving the performance of NISQ computers to execute existing quantum algorithms. One strategy is to mitigate errors through active or passive modifications to the quantum state preparation and measurement protocols. For example, readout errors can be mitigated through post-processing steps [16–35] and gate errors can be mitigated by systematically enlarging errors before extrapolating to zero error [36–41]. A complementary strategy to error mitigation is circuit optimization, also known as circuit compilation. In particular, there is no unique way to encode a quantum algorithm into a set of gates, and certain realizations of an algorithm may be better-suited for a given quantum device. One widely used tool is $t|ket\rangle$ [42], which contains a variety of architecture-agnostic and architecture-specific routines. For example, Clifford identities such

as $\text{CNOT}^2 = \text{Identity}$ are automatically recognized. There are also a variety of other toolkits for circuit optimization, including hardware-specific packages for quantum circuits [43–61]. Since `t|ket` is a generic framework that contains many algorithms that have already been benchmarked against other procedures, it will serve as our baseline.

We introduce two techniques that can be used to optimize circuits and that are complementary to existing methods. The first focuses on the identification of recurring sets of quantum gates in a circuit. Identifying such recurring sets of gates (RSG) can be very important, since any optimization of these RSGs has an enhanced effect on the overall circuit. Furthermore, identifying recurring gate sets can be useful for future hardware optimizations where the fidelity of certain common operations can be enhanced at the expense of other, less frequent operations. Depending on the operation(s), this optimization could be at the level of microwave pulse controls or it may require custom hardware architectures. The second technique optimizes a generic circuit by eliminating unnecessary gates or unused qubits such that the circuit depth becomes as short as possible. One example where such an optimization can lead to simplifications is a case where a quantum circuit has been designed with complete generality in mind. In this case, for a certain initial state the circuit only reaches a select set of intermediate states such that some operations become trivial and can be eliminated. The elimination of unnecessary gate operations introduced here focuses on controlled operations such as a Toffoli or a CNOT gate in a quantum circuit. The heart of the elimination technique resides in the identification of zero- or low-amplitude computational basis states, that allows us to determine whether the entire gate or (part of) qubit controls can be removed. Both of these techniques are combined in an optimization protocol called AQCEL (and pronounced “excel”) for *Advancing Quantum Circuit by ICEPP and LBNL*.

To demonstrate the effectiveness of these techniques, we will use a quantum algorithm from HEP to perform a calculation in quantum field theory. The particular algorithm that we study models a *parton shower*, which is the collinear final state radiation from energetic charged (under any force) particles [15]. This algorithm is a use-

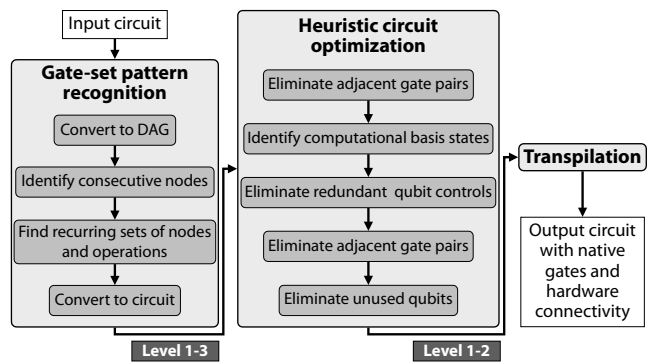


Figure 1: Flowchart of the proposed optimization protocol. The first part is the RSG pattern recognition, in which the circuit is converted into a directed acyclic graph (DAG) to identify recurring quantum gates. In the second part, we eliminate unnecessary gates and unused qubits through a heuristic approach. Finally, the resulting circuit can be encoded into particular gates for specific hardware.

ful benchmark because it provides an exponential speedup over the most efficient known classical algorithm and the circuit depth can be tuned for precision. While we show results for this specific circuit, the proposed protocol has a wide range of applicability for quantum computing applications across science and industry.

This paper is organized as follows. Section 2 provides an overview of the AQCEL protocol. The application of this protocol to the HEP example is presented in Sec. 3. Following a brief discussion in Sec. 4, the paper concludes in Sec. 5.

2 AQCEL optimization protocol

As already mentioned, the AQCEL protocol comprises two components: identifying recurring quantum gates (Sec. 2.1) and eliminating unnecessary gates and unused qubits (Sec. 2.2). This approach focuses on circuit optimization at the algorithmic level instead of at the level of a specific implementation using the native gates for a particular quantum device. A high-level flowchart for our protocol is presented in Fig. 1. The individual optimization steps are described below.

2.1 Gate set pattern recognition

First, the AQCEL attempts to identify gate set patterns in an arbitrary quantum circuit and extract RSGs from the circuit.




Figure 2: An example circuit containing a Toffoli gate (top) and its corresponding DAG (bottom).

2.1.1 Representation in directed acyclic graph

In a quantum circuit, individual qubits are manipulated by gate operations one by one, meaning that the quantum state represented at a certain point of the circuit should not be affected by gate operations applied afterward (at a later point in the circuit). Such a structure can be described by a directed acyclic graph (DAG). A DAG allows us to easily check dependencies between qubits and to extract a subset of the circuit that functions for certain tasks.

First, we convert a quantum circuit to the form of a DAG using the DAGCircuit class in Qiskit Terra API, where a node represents an operation by a quantum gate and an edge that connects the nodes represents a qubit. In the case of a Toffoli gate, the node corresponding to the Toffoli gate has three incoming edges (qubits before the gate operation) and three outgoing edges (qubits after the gate operation). Figure 2 shows an example circuit containing a Toffoli gate and its corresponding DAG.

A gate set pattern recognition for a quantum circuit can be resolved through the DAG

representation. The identity of the RSG functionality can be ensured by checking the identity of DAGs of two circuits, as a graph isomorphism problem. The algorithm of gate set pattern recognition consists of two steps: (1) finding RSG candidates with DAG representation, which enables to access recurring gates efficiently and identify a dependency of two gates, using depth-first search with heuristic pruning, and (2) checking the DAG isomorphism by graph hashing with Weisfeiler Lehman graph hash [62] implemented in the NetworkX library [63]. The details of the gate set pattern recognition including computational complexity are given in Appendix A, with the pseudocode of the algorithm.

2.1.2 Tiered extraction of recurring gate sets

The appearance pattern of RSGs in a quantum circuit may depend on specific encoding of the quantum algorithm. To account for different patterns, we consider three different levels of matching criteria to define the recurrence of quantum gates:

Level 1 : Only matching in gate types,

Level 2 : Matching in gate types and the *roles* of qubits that the gates act on,

Level 3 : Matching in gate types and both *roles* and *indices* of qubits that the gates act on.

The matching criterion in Level 1 is the least stringent: it just identifies the same sets of quantum gates appearing in the circuit, irrespective of which qubits they act on. The Level 2 is more strict and ensures that the qubits the RSGs act on have the same roles. In other words, the qubit connections between the gates inside a single RSG are maintained but the qubit indices might vary between the RSGs. The Level 3 applies the most stringent condition, where the qubits that the RSGs act on must have the same roles and qubit indices. In other words, the RSGs must appear on an identical set of qubits in the circuit. The appearance patterns of the RSGs are illustrated in Fig. 3 for the three matching criteria.

The identified RSGs are ranked in terms of the product of the number of gates constituting the set and the number of occurrence of the set in the circuit. A fixed number of top-ranked RSGs are extracted from the circuit in this step.

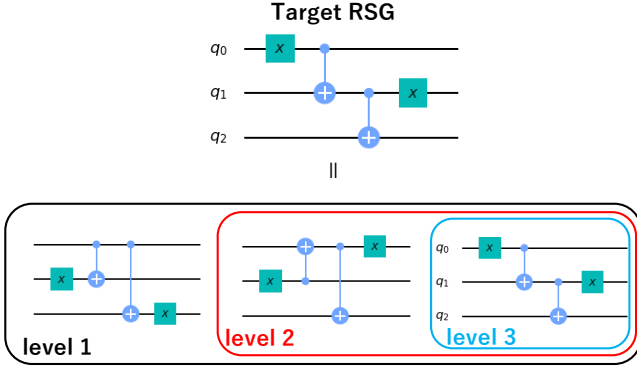


Figure 3: Possible RSG patterns for a given target RSG corresponding to the three levels of matching criteria.

2.2 Heuristic circuit optimization

After attempting to identify RSGs in the circuit, a heuristic optimization procedure takes place to make the circuit depth as short as possible by eliminating redundant gates or unused qubits. In this step, we consider two levels of optimization:

Level 1 : Optimize the entire circuit including RSGs,

Level 2 : Optimize the entire circuit, but for the RSGs only adjacent gate pairs (see below) are removed.

The Level 1 optimization would provide a shorter, more efficient circuit. Compared to Level 1, Level 2 optimization likely results in a deeper circuit for most cases, while it provides more room for improvement in later compilation stages if the RSGs have specialized low-level implementations.

2.2.1 Basic idea of redundant controlled operations removal

A controlled operation with such as a CNOT or a Toffoli gate performs a different operation depending on the quantum state of the system at the point where the gate is applied. Let m be the number of control qubits of this operation. Consider expanding the state of the full system $|\psi\rangle$ into a superposition of computational basis states as

$$|\psi\rangle = \sum_{j,k} c_{j,k} |j\rangle_{\text{ctl}} \otimes |k\rangle, \quad (1)$$

where $|\cdot\rangle_{\text{ctl}}$ denotes the state of the control qubits, while the unsubscripted ket corresponds to the rest of the system. We write the states as integers with $0 < j < 2^m - 1$ and $0 < k < 2^{n-m} - 1$. We assume that the controlled operation for the

gate is applied when all control qubits are in the $|1\rangle$ state, which corresponds to the state $|j\rangle_{\text{ctl}} = |11\dots 1\rangle = |2^m - 1\rangle_{\text{ctl}}$. This allows to classify the state of the system into three general classes using the amplitudes $c_{j,k}$:


Triggering : $c_{j,k} \neq 0$ iff $j = 2^m - 1$. The controlled operation of the gate in question is applied for all computational bases in the superposition.

Non-triggering : $c_{2^m-1,k} = 0$ for all k . The controlled operation is never applied.

Undetermined : The state is neither triggering nor non-triggering.

A circuit containing triggering or non-triggering controlled gates can be simplified by removing all controls (triggering case) or by eliminating the gates entirely (non-triggering case). While an undetermined single-qubit controlled gate cannot be simplified under the current scheme, an undetermined multi-qubit controlled gate can be by removing the controls on some of the qubits, if the state of the system satisfies the condition described in Appendix C.

As an example of this concept, consider the following simple circuit:



If the second qubit is in the initial state $|0\rangle$, the first CNOT gate has no effect and can be removed from the circuit as the $|0\rangle$ is the non-triggering state of CNOT. The second qubit before the second CNOT gate is in the state $|1\rangle$, which is the triggering state. Therefore, the qubit control can be removed from the second CNOT gate. The first two qubits before the Toffoli gate are in the superposition of $|01\rangle$ and $|11\rangle$, which is an undetermined state for the Toffoli gate. Since the Toffoli gate has a triggering bitstring $\{11\}$, and the second qubit is always in the $|1\rangle$ state, this second qubit control can be removed from the Toffoli gate, replacing it with a CNOT gate controlled only on the first qubit.

The heuristic circuit optimization therefore requires, for each controlled gate, the identification of possible states the control qubits can take, and

the removal of unnecessary parts of the controlled operations. These two steps are discussed in detail in the following.

It is well known that an arbitrary multi-qubit controlled- U gate with m control qubits can be decomposed into $\mathcal{O}(m)$ Toffoli and controlled- U gates [64]. Therefore, in the remainder of this paper, we assume that all controlled gates are reduced to Toffoli gates denoted as $C^2[X]$, and singly-controlled unitary operation denoted as $C[U]$. This implies that the only triggering bitstrings we need to consider are either $\{1\}$ or $\{11\}$. For a n -qubit circuit composed of N multi-qubit controlled- U gates, each having at most n control qubits, this decomposition results in at most

$$\tilde{N} = nN \quad (2)$$

controlled gates.

2.2.2 Identification of computational basis states

In general, a circuit consisting of n qubits creates a quantum state described by a superposition of all of the 2^n computational basis states. However, it is rather common that a specific circuit produces a quantum state where only a subset of the computational basis states has nonzero amplitudes. Moreover, the number of finite-amplitude basis states depends on the initial state. This is why the three classes of the states of the system arise.

The state classification at each controlled gate can be determined either through a classical simulation or by measuring the control qubits repeatedly. In the case of a classical simulation, one can either perform the full calculation of the amplitudes, or simply track all the computational basis states whose amplitudes may be nonzero at each point of the circuit without the calculation of the amplitudes. AQCEL adopts the latter method for the lower computational resource. When instead the quantum measurements are used, the circuit is truncated right before the controlled gate in question, and the control qubits are measured repeatedly at the truncation point. Finiteness of the relevant amplitudes can be inferred from the distribution of the obtained bitstrings, albeit within the statistical uncertainty of the measurements.

A few notes should be taken on the computational costs of the two methods. Consider an

n -qubit circuit with N controlled gates. As discussed before, reducing this to either $C^2[X]$ or $C[U]$ results in $\mathcal{O}(\tilde{N})$ single or double controlled gates. A classical simulation of the state vector before a given controlled gate has an exponential scaling in the number of qubits and requires $\mathcal{O}(2^n)$ computations. On the other hand, measuring the $m = 1$ or 2 control qubits M times, which results in M bitstrings of length m , only requires $\mathcal{O}(M)$ operations. Repeating this for all \tilde{N} gates requires $\mathcal{O}(\tilde{N}2^n)$ for the classical simulation and $\mathcal{O}(\tilde{N}^2M)$ when using quantum measurements.

More details on the estimates of the computational resource necessary for the identification of computational basis states, as well as other optimization steps, are described in Appendix B.

2.2.3 Elimination of redundant controlled operations

Once the nonzero-amplitude computational basis states are identified at each controlled gate, we remove the gate or its controls if possible. When using classical simulation, the entire circuit is analyzed first before the control elimination step. When quantum measurements are instead used, circuit execution, measurements, and circuit optimization are performed at each controlled gate separately.

The control elimination step for each controlled gate proceeds as follows. For a $C[U]$ gate, compute the probability of observing $|1\rangle$ of the control qubit. If that probability is 1, eliminate the control and only keep the single unitary gate U . If the probability is 0, remove the controlled gate from the circuit. In all other cases, keep the controlled gate. For a $C^2[X]$ (Toffoli) gate, compute the probabilities of the four possible states $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$. If the probability of $|11\rangle$ is 1, remove the two controls and only keep the X gate. If the probability of $|11\rangle$ is 0, remove the entire Toffoli gate. If neither of those two conditions are true (the undetermined class), it is still possible to eliminate one of the two controls. This is true if the probability of the state $|01\rangle$ ($|10\rangle$) is zero, in which case one can eliminate the first (second) control. The following pseudocode is the full algorithm for redundant controlled operations removal.

Algorithm 1: Redundant controlled operations removal

```

for all  $C[U]$  or  $C^2[X]$  gate  $g$  in the circuit do
  execute circuit up to, but not including,  $g$ 
  if  $g$  is a  $C[U]$  gate then
    measure the control qubit  $q$  in the  $Z$  basis
    multiple times
    if  $\{1\}$  is observed in the measurement re-
    sults then
      if  $\{0\}$  is not observed in the measure-
      ment results then
        turn  $g$  into a  $U$  gate acting on the
        target qubit
      end if
    else
      eliminate  $g$ 
    end if
  else
    measure the control qubits  $q_1q_2$  in the  $Z$ 
    basis multiple times
    if  $\{11\}$  is observed in the measurement re-
    sults then
      if  $\{00\}$  and  $\{01\}$  and  $\{10\}$  are not ob-
      served in the measurement results then
        turn  $g$  into an  $X$  gate acting on the
        target qubit
      else if  $\{01\}$  is not observed in the mea-
      surement results then
        eliminate the control on  $q_1$ 
      else if  $\{10\}$  is not observed in the mea-
      surement results then
        eliminate the control on  $q_2$ 
      end if
    else
      eliminate  $g$ 
    end if
  end if
end for

```

Note that for noisy quantum circuits the measurements of the states will not be exact, and one expects errors in the probabilities to observe certain bitstrings. This means that one has to impose thresholds when deciding whether we call the state triggering, non-triggering or undetermined. Once such a threshold has been decided, the number of measurements required has to be large enough for the statistical uncertainty to be smaller than this threshold. This will be discussed in more detail in Sec. 3 when we give explicit examples.

The computational cost of determining whether to eliminate controls or the entire

controlled operation can easily be determined. Given the measured bitstrings, which as discussed in the previous section can be determined with $\mathcal{O}(\tilde{N}^2M)$ operations, one can compute the probabilities for each possible bitstring, and therefore decide whether to simplify a controlled operation using $\mathcal{O}(\tilde{N})$ operations. Some more details about the resource scaling are given in Appendix B.

Note that superfluous controlled operations can also be found and eliminated using the so-called ZX calculus [65, 66]. In fact, the ZX calculus is complete in the formal logic sense of the word, such that one can always prove that an unnecessary gate can be removed using the ZX calculus. However, in general this requires exponential resources, and therefore has no scaling advantage with respect to simply computing the state vectors. Of course, the ZX calculus is still incredibly powerful and underlies many of the optimization techniques of quantum transpilers, such as the t|ket> compiler we compare to later.

2.2.4 Elimination of adjacent gate pairs

Note that if a unitary operator A and its Hermitian conjugate A^\dagger act on the same set of qubits adjacently, resulting in an identity operation, the gates implementing these operators can be removed from the circuit. While this is an obvious simplification, the removal of gates through the optimization steps described above can result in a circuit with such cancelling gate pairs. For this reason, this step of gate reduction is applied before and after eliminating redundant controlled operations.

2.2.5 Elimination of unused qubits

After taking the above steps, the circuit is examined for qubits where no gate is applied (identically remain in the initial state), which are then removed from the circuit. Such a situation occurs e.g., when a quantum circuit designed to work universally with different initial states is executed with a specific initial state. An example of such a circuit is the sequential algorithm we consider in the next section.

3 Application to quantum algorithm

The circuit optimization protocol described in Sec. 2 has been deployed to a quantum algorithm designed for HEP [15]. The heuristic optimization (Sec. 2.2) is performed at Level 1 for the optimization on existing quantum hardware. In our results, we present how many gates are removed in three steps of the heuristic optimization namely

- Step 1: Eliminate adjacent gate pairs
- Step 2: Eliminate redundant controlled operations
- Step3: Eliminate adjacent gate pairs again

3.1 Quantum parton shower algorithm

Simulating quantum field theories is a flagship scientific application of quantum computing. It has been shown that a generic scattering process can be efficiently simulated on a quantum computer with polynomial resources [67]. However, such circuits require prohibitive resources in the context of near-term devices.

A complementary approach is to simulate one component of the scattering process. In particular, Ref. [15] proposed an algorithm to simulate the collinear radiation from particles that carry a nonzero fundamental charge. Such radiation approximately factorizes from the rest of the scattering amplitude and can therefore be treated independently. This factorization is the basis for parton shower Monte Carlo generators in HEP. The quantum parton shower (QPS) algorithm provides an exponential speedup over known algorithms when the charge is not the same for all particles that can radiate.

The particular example demonstrated in Ref. [15] starts with n fermions that can be either type f_1 or f_2 . These fermions can radiate a scalar particle ϕ , which itself can split into fermion-anti-fermion pairs (of the same or different type). The relevant parameters are the three couplings g_1 , g_2 , and g_{12} between f_1 and ϕ , f_2 and ϕ , and $f_1 f_2$ ($f_1 f_2$) and ϕ , respectively. The shower evolution is discretized into N_{evol} steps and at each step, one of the particles could radiate / split or nothing happens. This produces a precise result when N_{evol} is large. Figure 4 shows the quantum circuit block for the m^{th} step of the quantum circuit.

First, the fermions are rotated into a new basis f_a and f_b where the effective mixing g_{ab} between $f_a f_b$ ($f_a f_b$) and ϕ is zero. Then, the number of particles of each type are counted and stored in registers n_a , n_b , and n_ϕ . Next, a Sudakov factor is calculated to determine if an emission happens (or not). This operation depends only on the total number of particles of each type. After the emission step, the particle and history registers are modified depending on the emission. Lastly, the fermions are rotated back into the f_1 and f_2 basis. Some of the steps in this algorithm are universal (independent of m) and some dependent on m due to the running of coupling constants with energy scale.

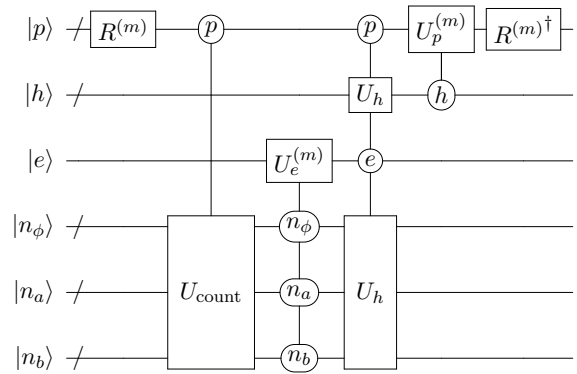


Figure 4: The m^{th} step of the quantum circuit for the algorithm proposed in Ref. [15]. There are three physical registers: $|p\rangle$ containing the set of particles at this step; $|h\rangle$ for the branching history; and $|e\rangle$ which is a binary variable representing the presence or absence of an emission at this step. The three lower registers count the number of particles of type ϕ , a , and b and are uncomputed before the end of the circuit. The exact form of the rotation matrices $R^{(m)}$ and the unitary operations U_{count} , $U_e^{(m)}$, U_h , and $U_p^{(m)}$ can be found in Ref. [15].

3.2 Experimental setup

The QPS simulation is implemented into a quantum circuit using IBM Qiskit version 0.22.0 [48] with Terra 0.15.2, Aer 0.6.1 and Ignis 0.4.0 APIs in Python 3.8 [68]. First, we attempt to optimize the circuits running on a classical computer with a single 2.4 GHz Intel core i5 processor.

In order to evaluate the AQCEL performance, the same QPS circuit optimized using t|ket> [42] in pytket 0.6.1 before transpilation is used as a reference. The optimization using t|ket> is done as follows. We consider the list of ten pre-defined passes described in [69]. The passes are tried one




Figure 5: Quantum circuit for the QPS simulation with two branching steps. The two RSGs, identified by the gate set pattern recognition step in the optimization scheme, are indicated by the ovals with different background colors.

by one on the QPS circuit, and the one that reduces the number of gates the most is applied to the circuit. The same set of passes are tried again on the resulting circuit to identify and apply the pass that most effectively reduces the gate count. This iterative process is repeated until the gate count is no longer reduced by any of the passes. The selected sequence of passes are used for evaluating the $t|\text{ket}\rangle$ performance in the remainder of the studies.

The QPS algorithm is executed on the 27-qubit “ibmq_sydney” [70], one of the IBM Quantum Falcon Processors, and the statevector simulator in Qiskit Aer with and without optimizing the circuit. For the results obtained solely from the statevector simulator, all the qubits are assumed to be connected to each other (referred to as the ideal topology). When executing the algorithm on the sydney, the gates in the circuit are transformed into machine-native single- and two-qubit gates, and the qubits are mapped to the hardware accounting for the actual qubit connectivity. For all the circuits tested with the sydney below, the noise-adaptive mapping is performed by taking into account the read-out and CNOT gate errors from the calibration data as well as the qubit

connection constraints¹. Gate cancellations also take place at this stage using the commutativity of native gates and unitary synthesis, as documented in Qiskit Terra API. This qubit mapping and gate cancellation process are repeated eleven times, and the circuit obtained with the smallest number of gates is finally tested with the sydney.

3.3 Results

3.3.1 Circuit optimization for $N_{\text{evol}} = 2$ branching steps using classical simulation

Circuit optimization performance of AQCEL is evaluated for a quantum circuit of the QPS simulation with $N_{\text{evol}} = 2$ branching steps assuming an ideal topology. The simulation does not consider any effects from hardware noise. The initial state is chosen to be $|f_1\rangle$, and the coupling constants are set to $g_1 = 2$ and $g_2 = g_{12} = 1$. Both $f \rightarrow f'\phi$ and $\phi \rightarrow f\bar{f}$ processes are considered². The original circuit constructed using Qiskit is shown in Fig. 5.

¹This corresponds to the transpilation of level 3 pass manager, as implemented in Qiskit Terra.

²Ref. [15] noted that when these are unphysically removed, the circuit can be simulated efficiently classically (see also Ref. [14]).

First, the RSG pattern recognition is performed against the circuit. When the Level 2 RSG pattern recognition is applied, two RSGs are identified, as also shown in Fig. 5, with the requirements on the number of nodes in each RSG being between 5 and 7 and the number of repetitions being 4 or more. If the matching level is raised from Level 2 to 3, candidate patterns with smaller numbers of nodes or repetitions are generally found.

Next, the heuristic optimization (Sec. 2.2) is performed over the entire circuit at Level 1. This step consists of identifying nonzero-amplitude computational basis states, removing redundant controlled operations, removing adjacent cancelling gate pairs (performed twice), and removing unused qubits. Nonzero-amplitude computational basis states are identified through classical calculation.

After the algorithmic level circuit optimization, the quantum gates in the circuit are decomposed into single-qubit gates (U_1, U_2, U_3) and CNOT gates. Figure 6 shows the numbers of the single-qubit and CNOT gates, the sum of the two, and the depth of the circuit before and after the optimization. The circuit depth is defined as the length of the longest path from the input to the measurement gates, with each gate counted as a unit, as implemented in Qiskit. The figure compares the values from the original circuit, the circuit optimized with $t|ket\rangle$ only, that with AQCEL only, and that with the combination of the two. The AQCEL optimizer reduces the total number of gates by 54%, resulting in a 51% reduction of the circuit depth. In particular, the reduction of the number of CNOT gates is 47%. This compares to $t|ket\rangle$, which reduces the total number of gates by 26%, CNOT by 1%, and the circuit depth by 10%. This means that, for the QPS algorithm, AQCEL is 38% more efficient than $t|ket\rangle$ in reducing the gate counts, and 46% more specifically for CNOT, and makes the circuit 45% shorter. Combining the two optimizers, the gate count is reduced by 63% (50% for CNOT only) and the depth by 55% with respect to the original circuit. The combined optimizer is 51% more efficient than the $t|ket\rangle$ alone for gate reduction (49% for CNOT only), producing a 50% shorter circuit for the QPS algorithm.

For the AQCEL optimizer, the gate reduction occurs mostly at the stage where the redundant




Figure 6: Numbers of single-qubit ($U_{1,2,3}$) gates, CNOT gates and the sum of the two as well as the depth of the two-branching step QPS circuit decomposed into native gates before and after optimization. The computational basis states with nonzero amplitudes at controlled gates are identified using classical calculation in the heuristic optimization step of AQCEL.

qubit controls are removed. Starting with 1279 gates (excluding barrier and measurement gates), the first adjacent gate-pair elimination step, the redundant qubit control reduction step, and the second gate-pair elimination step each removes 170, 510 (40% of the 1279 gates), and 6 gates, respectively. In terms of the computational cost, the wall time is by far dominated by the two adjacent gate-pair elimination steps combined, accounting for 91% of the total time, with a subdominant contribution of 7% from the redundant qubit control reduction.

Finally, the number of qubits is reduced from 24 to 21 with the AQCEL optimizer, while it is unchanged by $t|ket\rangle$. One qubit is removed from each of the three registers n_a, n_b , and n_ϕ because those qubits are used only for $N_{\text{evol}} \geq 3$ branching steps.

3.3.2 Circuit optimization for $N_{\text{evol}} = 1$ branching step using classical simulation and quantum measurements

The quantum circuit for the two-branching step QPS simulation is still too deep to produce useful results on a real existing quantum computer, even after optimizing the circuit. Therefore, we consider the circuit with only one branching step using the sydney and the statevector simulator. The initial state, coupling constants, and considered processes are the same as those used for the

$N_{\text{evol}} = 2$ branching steps simulation.

First, we examine the gate and qubit counts for the one-branching step QPS simulation assuming an ideal topology. Starting with 486 gates, the AQCEL optimizer removes 24, 260 (53% of 486 gates), and 2 gates in the three steps of the heuristic optimization, in the order given above. The adjacent gate-pair elimination step still dominates the wall time (96%). However, the redundant qubit control reduction now takes about 9 times less time than that for the two-branching step simulation, which is consistent with the exponential behavior of the computing cost of the step, as discussed in Sec. 2. The number of qubits is reduced from 15 to 13 with the AQCEL optimizer. One of four ancilla qubits is removed because three ancillas are sufficient for decomposing all the multi-controlled gates in the $N_{\text{evol}} = 1$ step. The register n_ϕ , composed of only one qubit, is also removed because it is used only for the case where the initial state is $|\phi\rangle$.

Next, the optimized circuits are transpiled considering the qubit connectivity of the `ibmq_sydney`. Figure 7 shows the same set of distributions as in Fig. 6, but for the one-branching step QPS simulation with the sydney-specific transpilation. The AQCEL optimizer reduces the number of native gates significantly for the one branching step as well. The relative reduction is more drastic for the one branching step than the two branching steps, mainly because the former (shallow) circuit has relatively more zero-amplitude computational basis states than the latter (deep) circuit.

We then evaluate the performance of the optimizers using the `sydney`. A particular challenge when employing AQCEL with a real quantum computer is in the determination of the bitstring probabilities of the control qubits at each controlled gate using quantum measurements. Due to hardware noise, the list of observed bitstrings would contain contributions from errors on the preceding gates and the measurement itself.

To mitigate the measurement errors, we obtain the correction by measuring the calibration matrix for the control qubits (with 8192 shots per measurement) using Qiskit Ignis API. The correction is then applied to the observed distribution with a least-squares fitting approach.

The errors incurred by gate imperfection accumulate throughout the circuit execution and are




Figure 7: Numbers of single-qubit ($U_{1,2,3}$) gates, CNOT gates and the sum of the two as well as the depth of the one-branching step QPS circuit transpiled considering the `ibmq_sydney` topology before and after the optimizations. The computational basis states with nonzero amplitudes at controlled gates are identified using classical calculation in the heuristic optimization step of AQCEL.

difficult to correct. Instead, in AQCEL, we opt to ignore the observed bitstrings with occurrence below certain thresholds (cutoff thresholds), under the assumption that the gate errors act as a perturbation that inserts spurious computational basis states with small amplitudes into the superposition of the system. This can be improved in the future with additional computational complexity by using gate error mitigation such as the zero noise extrapolation mentioned in Sec. 1.

The cutoff thresholds are defined as follows. We consider the errors in single-qubit ($U_{1,2,3}$) and two-qubit (CNOT) gates separately for all the hardware qubits. The reported error rates at the time of the experiment, measured during the preceding calibration run of the hardware, are used for the calculations. Let the $U_{1,2,3}$ and CNOT error rates be $\epsilon_U^{(i)}$ and $\epsilon_{CX}^{(i,j)}$, respectively, with i and j indicating qubits that the gates act on. We can approximately calculate the probability p_ϵ of measuring the states with at least one gate error occurring anywhere in the circuit by performing qubit-wise (index-dependent) multiplications of the error rates:

$$p_\epsilon = 1 - \left[\prod_i \left(1 - \epsilon_U^{(i)}\right)^{n_U^{(i)}} \prod_{i \neq j} \left(1 - \epsilon_{CX}^{(i,j)}\right)^{n_{CX}^{(i,j)}} \right] \sim N_{CX} \epsilon_{CX}, \quad (3)$$

where $n_U^{(i)}$ and $n_{CX}^{(i,j)}$ are the number of $U_{1,2,3}$ and

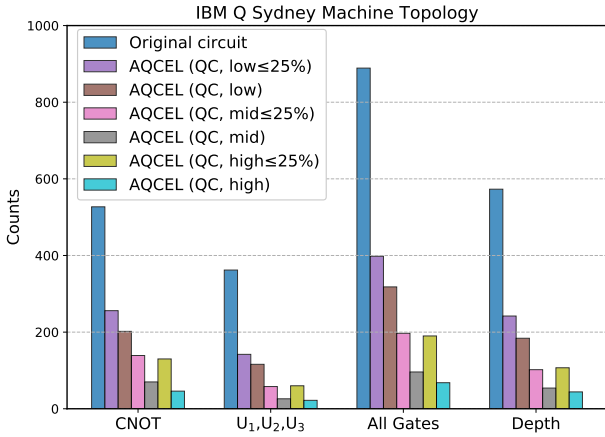


Figure 8: Numbers of single-qubit ($U_{1,2,3}$) gates, CNOT gates and the sum of the two as well as the depth of the one-branching step QPS circuit transpiled considering the `ibmq_sydney` topology before and after optimization. The probabilities of observing various bitstrings in the control qubits are measured using the `sydney` in the heuristic optimization step, and the three dynamic cutoff thresholds of s_ϵ^{low} , s_ϵ^{med} and s_ϵ^{high} are applied.

CNOT gates acting on the corresponding qubits, respectively, and in the last line we have assumed that all CNOT errors are equal, much larger than single gate errors and much smaller than one $\epsilon_U^{(i)} \ll \epsilon_{CX}^{(i,j)} = \epsilon_{CX} \ll 1$. The first cutoff threshold is

$$s_\epsilon^{\text{high}} := p_\epsilon, \quad (4)$$

corresponding to making an extreme assumption that any occurrence of a gate error during circuit execution results in a specific bitstring being observed at the measurement, and attempting to discard that bitstring. The second threshold,

$$s_\epsilon^{\text{low}} := p_\epsilon/2^m, \quad (5)$$

where m is the number of the measured control qubits, is related to another extreme assumption that the gate errors result in a uniform distribution of all possible bitstrings. The third and final threshold is the average of the above two:

$$s_\epsilon^{\text{med}} := (s_\epsilon^{\text{low}} + s_\epsilon^{\text{high}})/2. \quad (6)$$

It should be noted that p_ϵ increases as the circuit execution proceeds, as it is obtained by multiplying the error rates of all the preceding gates in the circuit. As an alternative strategy, we also examine static thresholds s_ϵ^f that are kept constant throughout the circuit, with values between 5% and 40%. We also consider capping the dynamic thresholds s_ϵ^{low} , s_ϵ^{med} , and s_ϵ^{high} at 25% (the reason behind the 25% will be given later).




Figure 9: Numbers of single-qubit ($U_{1,2,3}$) gates, CNOT gates and the sum of the two as well as the depth of the one-branching step QPS circuit transpiled considering the `ibmq_sydney` topology before and after optimization. The probabilities of observing various bitstrings in the control qubits are measured using the `sydney` in the heuristic optimization step, and the static cutoff thresholds of s_ϵ^f are applied.

Discarding all bitstrings with occurrence under certain thresholds obviously introduces errors of its own. For example, we observe that discarding bitstrings using unbounded s_ϵ^{high} as the threshold for the one-branching step QPS simulation circuit results in an elimination of most of the controlled gates in the later part of the circuit, rendering the circuit practically meaningless. Therefore, the actual threshold to be used with AQCEL should be chosen considering the tradeoff between the efficiency of the circuit optimization and the accuracy of the optimized circuit³.

Figure 8 shows the gate counts obtained from AQCEL optimizations using actual measurements on the `sydney` under dynamic cutoff thresholds. The gate counts decrease as the threshold is raised from s_ϵ^{low} to s_ϵ^{high} , as expected. Figure 9 shows the same distributions obtained with static thresholds. Almost no gate survives under the threshold of 40%, which likely implies a significant loss of accuracy of the computation result.

The number of qubits is reduced from 15 to 13 with the threshold of s_ϵ^{low} or s_ϵ^{med} , and to 11 with s_ϵ^{high} . Under static thresholds, the number of qubits is reduced from 15 to 13 for $5\% \leq s_\epsilon^f \leq 25\%$, but a significant reduction to 8 is seen for

³In the actual implementation, the threshold corresponding to 5% of the number of shots is applied to all the three cases to suppress contributions from imperfect measurement error mitigation.




Figure 10: Fidelity F_{meas} versus the number of native gates for the one-branching step QPS circuit transpiled considering the `ibmq_sydney` topology before and after optimization. The computational basis states with nonzero amplitudes at controlled gates are identified using classical calculation in the heuristic optimization step. These transpiled circuits are executed on the sydney to obtain the F_{meas} .

$s_{\epsilon}^f \geq 35\%$.

To evaluate the accuracy of the optimized circuit, we consider a classical fidelity of the final state of the circuit, which is defined in terms of the probability distribution of the computational basis states observed in the measurement at the end of the circuit. This quantity, denoted as F and referred to as just “fidelity” hereafter, is given by

$$F = \sum_k \sqrt{p_k^{\text{orig}} p_k^{\text{opt}}}, \quad (7)$$

where the index k runs over the computational basis states. The quantities p_k^{orig} and p_k^{opt} are the probabilities of observing k in the original and optimized circuits, respectively.

In fact, we compute two fidelity values for each optimization method. The first, denoted F_{sim} , aims to quantify the amount of modifications to the original circuit affected by the optimization at the algorithmic level. To calculate F_{sim} , both p^{orig} and p^{opt} are computed using the state vector simulation. The value of $F_{\text{sim}} = 1$ indicates that the optimized circuit is identical to the original circuit (up to a possible phase difference on each of the qubits), while a value different from unity gives a measure of how much the optimization has modified the circuit.

The second fidelity value, F_{meas} , is computed using measurements from actual quantum com-

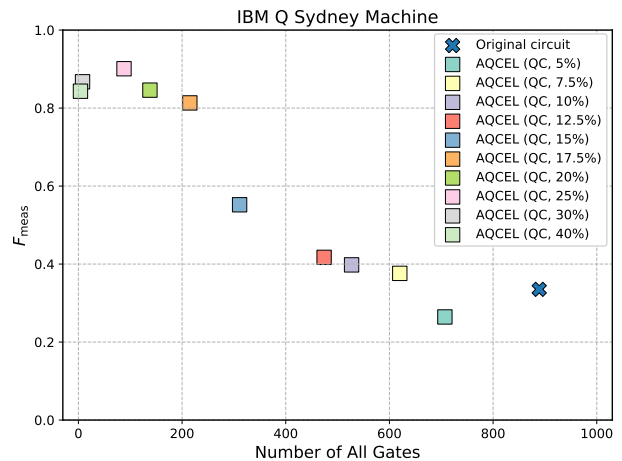


Figure 11: Fidelity F_{meas} versus the number of native gates for the one-branching step QPS circuit transpiled considering the `ibmq_sydney` topology before and after optimization. The probabilities of observing various bitstrings in the control qubits are measured using the sydney in the heuristic optimization step, and the static thresholds of s_{ϵ}^f are applied. These transpiled circuits are executed on the sydney to obtain the F_{meas} .

puter for p^{opt} . The rate of occurrence of a bitstring in a large number of repeated measurements is used as an estimate of p^{opt} . The p^{orig} is computed using simulation as for the F_{sim} . Even if the optimized circuit is identical to the original circuit, the presence of noise will mean $F_{\text{meas}} < 1$, with the discrepancy getting larger when more gates (in particular CNOT gates) are present in the circuit. Removing CNOT gates to obtain the optimized circuit will lower the overall effect of noise, which raises the value of F_{meas} . However, in some cases the CNOT gate removal also affects low-amplitude computational basis states, meaning the optimized circuit can differ from the original circuit, that might result in a suppression of the F_{meas} value. Thus, F_{meas} is a measure that takes into account the tradeoff of making the circuit shorter and changing the circuit through optimization.

Figure 10 shows the fidelity F_{meas} versus the number of native gates ($U_{1,2,3}$, CNOT) before and after optimization, where the optimization is performed using the classical simulation. One can see that shortening the circuit is increasing the F_{meas} as expected. The measurements are performed 81920 times for each of the circuit to obtain the F_{meas} values, and measurement error mitigation is not used in these or following F_{meas} measurements.

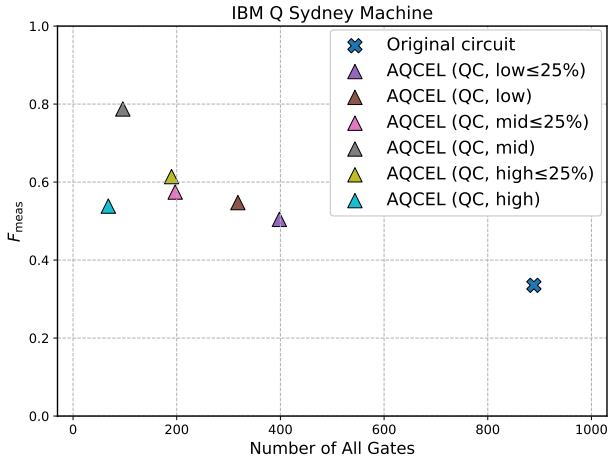


Figure 12: Fidelity F_{meas} versus the number of native gates for the one-branching step QPS circuit transpiled considering the `ibmq_sydney` topology before and after optimization. The probabilities of observing various bitstrings in the control qubits are measured using the `sydney` in the heuristic optimization step, and the three dynamic thresholds of $s_{\epsilon}^{\text{low}}$, $s_{\epsilon}^{\text{med}}$ and $s_{\epsilon}^{\text{high}}$ are applied. These transpiled circuits are executed on the `sydney` to obtain the F_{meas} .

When the elimination of redundant qubit controls is performed based on measurements using quantum computer and with static thresholds, the F_{meas} versus gate counts become those shown in Fig. 11 for various values of the s_{ϵ}^f thresholds. We observe that with increasing s_{ϵ}^f value the F_{meas} first increases, which indicates more aggressive optimization, but that at some point the F_{meas} starts to worsen, signaling that the optimized circuit becomes too far from the original circuit. For the circuit considered here, the performance of the optimization seems best for $s_{\epsilon}^f \sim 25\%$.

In Fig. 12 we show the result of the optimization with the dynamic thresholds of $s_{\epsilon}^{\text{high}}$, $s_{\epsilon}^{\text{med}}$, and $s_{\epsilon}^{\text{low}}$. We also show results for the capped variants, where the threshold is capped at 25%. The F_{meas} generally improves with higher thresholds, but is worse with $s_{\epsilon}^{\text{high}}$ than with $s_{\epsilon}^{\text{med}}$ because the accuracy is significantly degraded due to too aggressive threshold in the former case. The capped variants leave more gates in the circuit and have lower F_{meas} than the unbounded cases, except for when using $s_{\epsilon}^{\text{high}}$, where capping of the threshold seems to mitigate the loss of accuracy from overly aggressive optimization.

The results obtained from different approaches for finding nonzero-amplitude basis states and

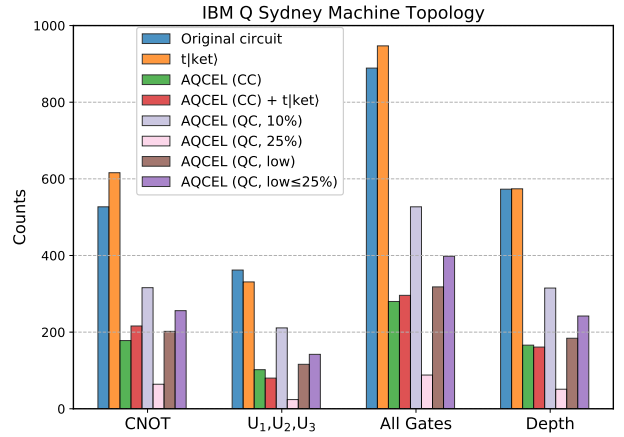


Figure 13: Numbers of single-qubit ($U_{1,2,3}$) gates, CNOT gates and the sum of the two as well as the depth of the one-branching step QPS circuit transpiled considering the `ibmq_sydney` topology before and after optimization under different schemes.

different choices of cutoff thresholds are summarized in Figs. 13 and 14 for comparison. It is worth noting that most of the AQCEL-based optimization shown in the figure improve the F_{meas} value over the `t|ket)`-only optimization. Another interesting finding is that the determination of bitstring probabilities with quantum measurements brings a better performance than the identification of nonzero amplitudes with classical calculation, if the cutoff threshold is set properly (25% for this case). A qualitative explanation for this would be that the quantum measurements and the cutoff serve to remove qubit controls over low-amplitude basis states, where such states contribute little to the final calculation result but the existence of the controlled gates produces those spurious states under the effect of hardware noise. An exact identification of computational basis states with nonzero amplitudes with classical simulation does not allow removing such qubit controls, effectively degrading the F_{meas} .

Figure 15 shows the fidelities F_{sim} versus F_{meas} before and after optimization under different schemes. The figure shows that the F_{sim} is identical to unity for all circuits optimized using the classical simulation, validating that the optimization has not affected the computational accuracy with respect to the original circuit. Although it also shows that the F_{sim} is slightly lowered from unity for all circuits optimized using actual measurements on the `ibmq_sydney`, the F_{meas} is bet-




Figure 14: Fidelity F_{meas} versus the number of native gates for the one-branching step QPS circuit transpiled considering the `ibmq_sydney` topology before and after optimization under different schemes. These transpiled circuits are executed on the sydney to obtain the F_{meas} .

ter than the original circuit because gate reductions suppress the effect of hardware noise.

4 Discussion

4.1 Applicability of proposed heuristic optimization

The core component of the proposed heuristic circuit optimization is the identification of computational basis states with nonzero amplitudes and the subsequent elimination of redundant controlled operations. Therefore, AQCEL is expected to work more efficiently for quantum algorithms in which the quantum state has a small number of high-amplitude computational basis states. In other words, if all the computational basis states have non-negligible amplitudes, AQCEL would not be effective. An example of when AQCEL is not effective is a quantum algorithm where an equal superposition state is first created by applying $H^{\otimes n}$ to the initial $|0\rangle^{\otimes n}$ state of the n -qubit system, such as Quantum Phase Estimation [71] and Grover's Algorithm.

4.2 Possibility of further simplifications

For certain quantum circuits, there is a case where there are successive multi-qubit controlled gates that share the same control qubits. One example is in the QPS simulation circuit (Fig. 4). The circuit determines if an emission happens




Figure 15: Fidelities F_{meas} versus F_{sim} for the one-branching step QPS circuit transpiled considering the `ibmq_sydney` topology before and after optimization under different schemes. These transpiled circuits are executed on the sydney to obtain the F_{meas} and a statevector simulator to obtain the F_{sim} . The vertical bars represent the total number of gates in the circuits (gauged by the right-hand vertical axis).

and which particle radiates or splits, depending on the total counts of particles of each type. These steps (corresponding to the blocks with controlled unitary operations denoted by $U_e^{(m)}$ and U_h in Fig. 4) require a lot of successive multiple controlled operations that share the same control qubits. In this case, if the circuit is expanded by adding an ancilla qubit and the triggering decision of the control qubits is stored into the ancilla qubit, the remaining multi-qubit controlled gates can be controlled by the ancilla. A potential caveat is that adding ancilla qubits might introduce additional SWAP gates when implementing the circuit to hardware. However, since this approach does not depend on the amplitudes of computational basis states of a given circuit state, it is complementary to the AQCEL optimization scheme and will open the possibility of reducing the overall gate counts further.

Another interesting possibility is that if a circuit turns out to contain only a small number of basis states, the circuit state can be represented using fewer qubits than the original ones. Given that this might require a completely new computational basis, this is left for future work.

4.3 Implication to hardware implementations of quantum circuits

The techniques introduced in the AQCEL protocol, i.e., identification of most-frequently-appearing sets of quantum gates as RSGs and the removal of redundant qubit control operations, have implications to hardware implementation of quantum circuits.

First, the RSGs would be a prioritized target for better mapping to quantum hardware. For the QPS algorithm, the RSGs contain multi-qubit controlled gates like the Toffoli gate, as shown in Fig. 5. In this case, these RSGs are further decomposed into collections of native single- and two-qubit gates. Therefore, the depth of the transpiled circuit depends significantly on which hardware qubits the decomposed RSG gates are mapped on to. If the transpilation algorithm accounts for the frequency of the occurrence of the RSGs, an improved qubit mapping can be created such that frequently-used controlled gates are applied on neighboring qubits with better connectivities on the quantum hardware.

In comparison between the AQCEL and $t|ket\rangle$ optimizers (e.g., Figs. 6 and 7), the $t|ket\rangle$ performance on the gate reduction turns out to be sub-optimal for the QPS algorithm. This is largely due to the lack of ability in $t|ket\rangle$ to remove redundant controlled operations through the identification of nonzero-amplitude computational basis states. However, in certain cases, the $t|ket\rangle$ -optimized circuit ends up with even more gates than the original circuit, as seen in Fig. 7 (note that the original and $t|ket\rangle$ -optimized circuits are both *optimized* using the noise-adaptive mapping and gate cancellation, see Sec. 3.2). The $t|ket\rangle$ optimizes a circuit assuming that all the qubits are connected to each other. This indicates that the circuit optimized with this assumption could result in more SWAP gates once the hardware connectivity is taken into account. On the other hand, when the $t|ket\rangle$ optimization is performed on a circuit already transpiled to specific hardware configuration, the gate counts appear to further reduce⁴. Even in this case, the AQCEL optimization performed before transpilation still produces a shorter circuit than the $t|ket\rangle$.

⁴This transpilation is performed with tools in Qiskit Terra. The $t|ket\rangle$ also allows its own circuit compilation to specific hardware, but it is not examined in this study.

This clearly indicates the importance of how well the gate errors are controlled and unnecessary gate operations are removed as much as possible. Moreover, if a circuit is mainly composed of Level 3 RSGs, as in the case of the QPS circuit used here, the hardware quality of control qubits of the RSGs will become crucial for the circuit simplification procedure proposed in the AQCEL protocol.

5 Conclusion and outlook

We have proposed a new protocol, called AQCEL, for analyzing quantum circuits to identify recurring sets of gates and remove redundant controlled operations. The heart of the redundant controlled operations removal resides in the identification of zero- or low-amplitude computational basis states. In particular, this procedure can be performed through measurements using a quantum computer in polynomial time, instead of classical calculation that scales exponentially with the number of qubits. Although removing qubit controls triggered in low-amplitude states will produce a circuit that is functionally distinct from the original, it is observed that this may be a desirable feature in some cases under the existence of hardware noise. If a quantum circuit contains recurring sets of quantum gates, those gates will be considered as candidates for further optimization in terms of both gate synthesis and hardware implementation. In the proposed protocol, the underlying technique to identify recurring gate sets is demonstrated, leading to the possibility of hardware-aware optimization of such gates including dedicated microwave pulse controls.

We have explored the AQCEL optimization scheme using the quantum parton shower simulation, a prototypical quantum algorithm for high-energy physics. For this algorithm, the proposed scheme shows a significant reduction in gate counts with respect to $t|ket\rangle$, which is one of the industry-standard optimization tools, while retaining the accuracy of the probability distributions of the final state.

This feature opens the possibilities to extend this optimization scheme further in future. We have considered several scenarios of the thresholds applied to the measured computational basis states to take into account the gate errors. The

measurement error is accounted for using the calibration matrix approach, and this can be improved by adapting the unfolding technique developed in Ref. [16] and related approaches that use fewer resources [22–24, 33, 72] or further mitigate the errors [35]. A substantial contribution to the gate errors originates from CNOT gates. There are a variety of approaches to mitigate these errors, including the zero noise extrapolation mentioned in Sec. 1. With such errors mitigated, we could potentially lower the thresholds. The threshold choice has a large impact to the accuracy of the measured probability distributions, as in Fig. 14, therefore the precise control of the measurement and gate errors is crucial for this approach.

Acknowledgements

We acknowledge the use of IBM Quantum Services for this work. The views expressed are those of the authors, and do not reflect the official policy or position of IBM or the IBM Quantum team.

CWB and BN are supported by the U.S. Department of Energy, Office of Science under contract DE-AC02-05CH11231. In particular, support comes from Quantum Information Science Enabled Discovery (QuantISED) for High Energy Physics (KA2401032).

We would like to thank Ross Duncan and Bert de Jong for useful discussions about the ZX Calculus.

A Algorithms of graph pattern recognition

The pattern recognition algorithm of recurring set of quantum gates (RSG) is described in Algorithm 2. This algorithm is based on depth-first search with heuristic pruning.

First, RSG candidates are built from seeding a quantum gate (node) by seeking possible combinations of RSGs that have descending connected quantum gates. A target node used as a seed, i.e., the beginning node, is selected with post-order traversal with a memorization technique to avoid a repeating calculation. The computational complexity of the algorithm is $O(N_{\text{nodes}}!)$ ⁵. Due

⁵The i -th node has $N_{\text{nodes}} - i$ RSG candidates in the worst case. Therefore, the computational complexity of




Figure 16: An example of quantum circuit (left) and its subgraph ($G' = \{A, B, C\}$) removed in our pattern recognition algorithm (right). A functionality of the corresponding circuit depends on the intermediate gate (D).

to a large number of combinations of recurring gates, the complexity is worse than the typical complexity of a classical computer, $O(n_{\text{qubits}}!)$ or $O(2^{n_{\text{qubits}}})$, because of $N_{\text{nodes}} = n_{\text{gates}} \geq n_{\text{qubits}}$ in most cases, and therefore it loses the benefit of quantum computer. To reduce the computational complexity, we prune the RSG candidates by requiring the length of the longest path, the minimum number and the maximum number of elements in RSG. The requirement of the minimum number of elements rejects a trivial RSG (e.g. $G = \{X\}$). The computational complexity reduces to $O(N_{\text{nodes}}^{N_{\text{thr}}})$ ⁶ where N_{thr} is a threshold value for the pruning, and the classical computer can calculate this in polynomial time when N_{thr} is fixed. However, this algorithm sometimes causes ill-defined RSGs, as shown in Fig. 16. The functionality of the quantum circuit from such an RSG depends on the intermediate gate that is not used in the RSG. These RSGs are rejected in this algorithm by requiring that there is no node, which is both a child and a parent nodes but not an element of the RSG ($\exists g_i, g_j \subseteq G', \{g_k | g_i \rightarrow g_k, g_k \rightarrow g_j\} \not\subseteq G'$).

After building the RSG candidates, they are grouped by graph isomorphism using the Weisfeiler Lehman graph hash. The use of graph hash does not ensure that two graphs are isomorphic, but the accuracy is sufficient for our use case. For the Level-1 matching criteria which consider only gate types, we assign the gate type as a node feature and assign nothing for an edge feature. For the Level-2 matching criteria which consider both gate types and qubit *roles*, we assign the gate type as a node feature and assign the target or control

the combination of the number of child-node's RSG candidates is $O(N_{\text{nodes}}!)$.

⁶We take N_{thr} RSG candidates from N_{nodes} nodes. Therefore, the computational complexity is $\binom{N_{\text{nodes}}}{N_{\text{thr}}} \approx N_{\text{nodes}}^{N_{\text{thr}}}$.

label as an edge feature. For the Level-3 matching criteria which consider gate types, qubit *roles* and *indices*, we assign the gate type as a node feature and assign the absolute qubit index as an edge feature.

Finally, the top-k RSGs are selected based on the frequency times the graph size.

Algorithm 2: Gate set pattern recognition with DAG

```

for all quantum gate (node) ( $g_i$ ) in the circuit ( $G$ ) do
  for all subset ( $G'$ ) beginning with the target node ( $g_i$ ) do
    if the longest path is longer than the threshold then
      continue
    end if
    if number of elements in subset is out of thresholds then
      continue
    end if
    if  $\exists g_i, g_j \subseteq G', \{g_k | g_i \rightarrow g_k, g_k \rightarrow g_j\} \not\subseteq G'$  then
      continue
    end if
     $G'$  is a RSG candidate.
  end for
end for
Make a set of RSGs ( $S(h) = \{G' | \text{hash}(G') = h\}$ )
Select top-k sets of RSGs ( $S(h)$ ) ordering by the frequency  $|S(h)|$  times RSG size ( $|G'|$ )

```

B Computational resources for the proposed optimization scheme

The computational cost needed to perform the proposed optimization scheme is evaluated here. We consider a quantum circuit that contains n qubits and N multi-qubit controlled gates, each acting on m control qubits and one target qubit.

The elimination of adjacent gate pairs proceeds, for each gate, by checking a pair-wise matching to the next gate until the end of the gate sequence. Since the gate could contain at most n qubits, the computational cost would be $\mathcal{O}(nN)$.

The next step in the optimization scheme is the identification of computational basis states. If we

use the classical calculation for simply tracking all the computational basis states whose amplitudes may be nonzero at each point of the circuit without the calculation of the amplitudes, it requires the computation of $\mathcal{O}(N2^n)$ states, so grows exponentially with n . This method allows the lower computational resource than a state-vector simulation though it neglects rare eliminations of redundant controlled operations. If we measure the control qubits at each controlled gate M times using a quantum computer, the total number of gate operations and measurements is given by $M\{m + (1 + m) + (2 + m) + \dots + (N - 1 + m)\} = \frac{1}{2}MN(N - 1) + mMN$. Therefore, the computational cost grows polynomially in $\mathcal{O}(MN^2 + mMN)$.

We next consider removing redundant qubit controls from a controlled gate with m control qubits. Using a quantum computer that measures the m control qubits M times, the measured number of computational basis states is M if $M < 2^m$, otherwise 2^m . For the classical calculation, the number of basis states is 2^m . Imagine that we choose an arbitrary combination among 2^m possible combinations of new qubit controls on the same controlled gate. If we want to know whether the chosen combination can act as the correct qubit control, we would need to check, for a given measurement done previously with quantum computer, if all the possible “unmeasured” computational basis states satisfy the chosen one or not. This requires $\mathcal{O}(m2^m)$. Since this has to be checked for all the measurements, the cost would be $\mathcal{O}(Mm2^m)$ if $M < 2^m$, otherwise $\mathcal{O}(m4^m)$. Therefore, the overall computation for the determination of redundant qubit controls would become $\mathcal{O}(Mm4^mN)$ or $\mathcal{O}(m8^mN)$ for N multi-qubit controlled gates, each having 2^m computational basis states. The classical calculation would require $\mathcal{O}(m8^mN)$ as well.

It is known that an arbitrary multi-qubit controlled- U gate with m control qubits can be decomposed into $\mathcal{O}(m)$ Toffoli and two-qubit controlled- U gates [64]. Therefore, if a controlled gate in the circuit is decomposed in this way, then above computational cost for the redundant qubit controls would become $\mathcal{O}(mN)$. With this decomposition, the total number of gate operations and measurement increases due to $\mathcal{O}(m)$ times more controlled gates. However, the computational cost for the identifica-

tion of computational basis states becomes only $\frac{1}{2}mMN(mN-1)+2mMN$, so still behaves polynomially in $\mathcal{O}(m^2MN^2)$ in case the quantum computer is used. For the classical calculation, the cost becomes $\mathcal{O}(mN2^n)$.

The final step of the optimization scheme is the elimination of unused qubits. This is just performed by checking qubits that all the gates in the circuit act on, therefore the computational cost is $\mathcal{O}(nN)$.

Given that a controlled gate has at most $n-1$ control qubits, the total computational cost for the entire optimization steps is $\mathcal{O}(n^2MN^2)$ and $\mathcal{O}(nN2^n)$ if the computational basis state measurement is performed using a quantum computer and classical calculation, respectively.

C General conditions to eliminate qubit controls

Given a multi-qubit controlled gate $C^m[U]$ and a system in the “undetermined” state $|\psi\rangle$ following the classification in Section 2.2.1, we can derive the condition for removal of a part of the controls to be allowed in the following way.

Let x be the number of controls to be removed. Without loss of generality, the decomposition of $|\psi\rangle$ can be rewritten as

$$|\psi\rangle = \sum_{i,l,k} \tilde{c}_{i,l,k} |i\rangle_{\text{ctl}'} \otimes |l\rangle_{\text{free}} \otimes |k\rangle, \quad (8)$$

where $|\cdot\rangle_{\text{ctl}'}$ and $|\cdot\rangle_{\text{free}}$ are the states of the $m-x$ remaining control qubits and the x qubits from which the controls are removed. From Eq. (1),

$$|i\rangle_{\text{ctl}' } \otimes |l\rangle_{\text{free}} = |2^x i + l\rangle_{\text{ctl}'}, \quad (9)$$

and therefore

$$\tilde{c}_{i,l,k} = c_{2^x i + l, k}. \quad (10)$$

Applying the original controlled gate to $|\psi\rangle$ yields

$$C^m[U]|\psi\rangle = \sum_{j=0}^{2^m-2} \sum_k c_{j,k} |j\rangle |k\rangle + \sum_k c_{2^m-1,k} |2^m-1\rangle U|k\rangle, \quad (11)$$

where ket subscripts and the tensor product symbols are omitted for simplicity. In contrast, the

new gate with fewer controls give

$$C^{m-x}[U]|\psi\rangle = \sum_{i=0}^{2^{m-x}-2} \sum_{l,k} \tilde{c}_{i,l,k} |i\rangle |l\rangle |k\rangle + \sum_{l,k} \tilde{c}_{2^{m-x}-1,l,k} |2^{m-x}-1\rangle |l\rangle U|k\rangle. \quad (12)$$

For the removal of x qubit controls to be allowed, the right hand sides of Eqs. (11) and (12) must be identical. This requires

$$\sum_{l=0}^{2^x-2} \sum_k \tilde{c}_{2^{m-x}-1,l,k} |2^{m-x}-1\rangle |l\rangle U|k\rangle = \sum_{l=0}^{2^x-2} \sum_k c_{2^m-2^x+l,k} |2^m-2^x+l\rangle |k\rangle. \quad (13)$$

Denoting

$$U|k\rangle = \sum_{k'} u_{kk'} |k'\rangle \quad (14)$$

and recalling Eqs. (9), Eq. (13) implies (replacing $k' \leftrightarrow k$ on the left hand sides)

$$\sum_{l=0}^{2^x-2} \sum_{k,k'} \tilde{c}_{2^{m-x}-1,l,k'} u_{k'k} |2^{m-x}-1\rangle |l\rangle |k\rangle = \sum_{l=0}^{2^x-2} \sum_k c_{2^m-2^x+l,k} |2^{m-x}-1\rangle |l\rangle |k\rangle. \quad (15)$$

Then, with Eq. (10), we have

$$\sum_{k'} \tilde{c}_{2^{m-x}-1,l,k'} u_{k'k} = \tilde{c}_{2^{m-x}-1,l,k} \quad \forall l, k. \quad (16)$$

Equation (16) holds if the row vector $\{\tilde{c}_{2^{m-x}-1,l,k}\}_k$ is an eigenvector of the matrix u with eigenvalue 1 under right multiplication for all l , or if $\tilde{c}_{2^{m-x}-1,l,k} = 0$ for all l and k .

Since the cost of exactly computing the complex amplitudes of the quantum state is high, in AQCEL we only consider this second condition of $\tilde{c}_{2^{m-x}-1,l,k} = 0$. When using quantum measurements to estimate the bitstring probabilities at the control qubits, this requirement translates to observing no bitstring with 1 in the qubits that are considered for control removal.

References

- [1] John Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, Aug 2018.

- [2] Alex Mott, Joshua Job, Jean Roch Vlimant, Daniel Lidar, and Maria Spiropulu. Solving a Higgs optimization problem with quantum annealing for machine learning. *Nature*, 550(7676):375–379, 2017.
- [3] Alexander Zlokapa, Alex Mott, Joshua Job, Jean-Roch Vlimant, Daniel Lidar, and Maria Spiropulu. Quantum adiabatic machine learning with zooming. 2019.
- [4] Jay Chan, Wen Guan, Shaojun Sun, Alex Zeng Wang, Sau Lan Wu, Chen Zhou, Miron Livny, Federico Carminati, and Alberto Di Meglio. Application of Quantum Machine Learning to High Energy Physics Analysis at LHC using IBM Quantum Computer Simulators and IBM Quantum Computer Hardware. *PoS, LeptonPhoton2019:049*, 2019.
- [5] Koji Terashi, Michiru Kaneda, Tomoe Kishimoto, Masahiko Saito, Ryu Sawada, and Junichi Tanaka. Event Classification with Quantum Machine Learning in High-Energy Physics. *Comput. Softw. Big Sci.*, 5(1):2, 2020.
- [6] Wen Guan, Gabriel Perdue, Arthur Pesah, Maria Schuld, Koji Terashi, Sofia Vallecorsa, and Jean-Roch Vlimant. Quantum Machine Learning in High Energy Physics. *Machine Learning: Science and Technology*, 5 2020.
- [7] Alexander Zlokapa, Abhishek Anand, Jean-Roch Vlimant, Javier M. Duarte, Joshua Job, Daniel Lidar, and Maria Spiropulu. Charged particle tracking with quantum annealing-inspired optimization. 8 2019.
- [8] Cenk Tüysüz, Federico Carminati, Bilge Demirköz, Daniel Dobos, Fabio Fracas, Kristiane Novotny, Karolos Potamianos, Sofia Vallecorsa, and Jean-Roch Vlimant. Particle Track Reconstruction with Quantum Algorithms. volume 245, page 09013, 3 2020.
- [9] Ilya Shapoval and Paolo Calafiura. Quantum Associative Memory in HEP Track Pattern Recognition. *EPJ Web Conf.*, 214:01012, 2019.
- [10] Frederic Bapst, Wahid Bhimji, Paolo Calafiura, Heather Gray, Wim Lavrijsen, and Lucy Linder. A Pattern Recognition Algorithm for Quantum Annealers. *Comput. Softw. Big Sci.*, 4:1, 2019.
- [11] Annie Y. Wei, Preksha Naik, Aram W. Harrow, and Jesse Thaler. Quantum algorithms for jet clustering. *Phys. Rev. D*, 101(9):094015, 2020.
- [12] Souvik Das, Andrew J. Wildridge, Sachin B. Vaidya, and Andreas Jung. Track clustering with a quantum annealer for primary vertex reconstruction at hadron colliders. 2019.
- [13] Kyle Cormier, Riccardo Di Sipio, and Peter Wittek. Unfolding measurement distributions via quantum annealing. *JHEP*, 11:128, 2019.
- [14] Davide Provasoli, Benjamin Nachman, Christian Bauer, and Wibe A de Jong. A quantum algorithm to efficiently sample from interfering binary trees. *Quantum Science and Technology*, 5(3):035004, 2020.
- [15] Christian W. Bauer, Benjamin Nachman, Davide Provasoli, and Wibe A. De Jong. A quantum algorithm for high energy physics simulations. 2019.
- [16] Christian W. Bauer, Wibe A. De Jong, Benjamin Nachman, and Miroslav Urbanek. Unfolding quantum computer readout noise. *npj Quantum Inf.*, 6:84, 2020.
- [17] R. C. Bialczak, M. Ansmann, M. Hofheinz, E. Lucero, M. Neeley, A. D. O’Connell, D. Sank, H. Wang, J. Wenner, M. Steffen, A. N. Cleland, and J. M. Martinis. Quantum process tomography of a universal entangling gate implemented with Josephson phase qubits. *Nature Physics*, 6(6):409–413, June 2010.
- [18] Matthew Neeley, Radoslaw C. Bialczak, M. Lenander, E. Lucero, Matteo Mariantoni, A. D. O’Connell, D. Sank, H. Wang, M. Weides, J. Wenner, Y. Yin, T. Yamamoto, A. N. Cleland, and John M. Martinis. Generation of three-qubit entangled states using superconducting phase qubits. *Nature*, 467(7315):570–573, September 2010.
- [19] A. Dewes, F. R. Ong, V. Schmitt, R. Lauro, N. Boulant, P. Bertet, D. Vion, and D. Esteve. Characterization of a Two-Transmon Processor with Individual Single-Shot Qubit Readout. *Physical Review Letters*, 108(5):057002, February 2012.
- [20] Easwar Magesan, Jay M. Gambetta, A. D. Córcoles, and Jerry M. Chow. Machine Learning for Discriminating Quantum Measurement Trajectories and Improving Readout. *Physical Review Letters*, 114(20):200501, May 2015.

- [21] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe. Demonstration of a small programmable quantum computer with atomic qubits. *Nature*, 536(7614):63–66, August 2016.
- [22] Chao Song, Kai Xu, Wuxin Liu, Chui-ping Yang, Shi-Biao Zheng, Hui Deng, Qiwei Xie, Keqiang Huang, Qiujiang Guo, Libo Zhang, Pengfei Zhang, Da Xu, Dongning Zheng, Xiaobo Zhu, H. Wang, Y.-A. Chen, C.-Y. Lu, Siyuan Han, and Jian-Wei Pan. 10-Qubit Entanglement and Parallel Logic Operations with a Superconducting Circuit. *Physical Review Letters*, 119(18):180511, November 2017.
- [23] Ming Gong, Ming-Cheng Chen, Yarui Zheng, Shiyu Wang, Chen Zha, Hui Deng, Zhiguang Yan, Hao Rong, Yulin Wu, Shaowei Li, Fusheng Chen, Youwei Zhao, Futian Liang, Jin Lin, Yu Xu, Cheng Guo, Lihua Sun, Anthony D. Castellano, Haohua Wang, Chengzhi Peng, Chao-Yang Lu, Xiaobo Zhu, and Jian-Wei Pan. Genuine 12-qubit entanglement on a superconducting quantum processor. *Physical Review Letters*, 122(11):110501, March 2019. arXiv: 1811.02292.
- [24] Ken X. Wei, Isaac Lauer, Srikanth Srinivasan, Neereja Sundaresan, Douglas T. McClure, David Toyli, David C. McKay, Jay M. Gambetta, and Sarah Sheldon. Verifying Multipartite Entangled GHZ States via Multiple Quantum Coherences. *Physical Review A*, 101(3):032343, March 2020. arXiv: 1905.05720.
- [25] Vojtech Havlicek, Antonio D. Córcoles, Kristan Temme, Aram W. Harrow, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta. Supervised learning with quantum enhanced feature spaces. *Nature*, 567(7747):209–212, March 2019. arXiv: 1804.11326.
- [26] Yanzhu Chen, Maziar Farahzad, Shinjae Yoo, and Tzu-Chieh Wei. Detector Tomography on IBM 5-qubit Quantum Computers and Mitigation of Imperfect Measurement. *Physical Review A*, 100(5):052315, November 2019. arXiv: 1904.11935.
- [27] Ming-Cheng Chen, Ming Gong, Xiao-Si Xu, Xiao Yuan, Jian-Wen Wang, Can Wang, Chong Ying, Jin Lin, Yu Xu, Yulin Wu, Shiyu Wang, Hui Deng, Futian Liang, Cheng-Zhi Peng, Simon C. Benjamin, Xiaobo Zhu, Chao-Yang Lu, and Jian-Wei Pan. Demonstration of Adiabatic Variational Quantum Computing with a Superconducting Quantum Coprocessor. *arXiv:1905.03150 [quant-ph]*, May 2019. arXiv: 1905.03150.
- [28] Filip B. Maciejewski, Zoltán Zimborás, and Michał Oszmaniec. Mitigation of readout noise in near-term quantum devices by classical post-processing based on detector tomography. *Quantum*, 4:257, April 2020. arXiv: 1907.08518.
- [29] Miroslav Urbanek, Benjamin Nachman, and Wibe Albert de Jong. Quantum error detection improves accuracy of chemical calculations on a quantum computer. *Physical Review A*, 102(2):022427, August 2020. arXiv: 1910.00129.
- [30] Benjamin Nachman, Miroslav Urbanek, Wibe A. de Jong, and Christian W. Bauer. Unfolding Quantum Computer Readout Noise. *arXiv:1910.01969 [physics, physics:quant-ph]*, May 2020. arXiv: 1910.01969.
- [31] Kathleen E. Hamilton and Raphael C. Pooser. Error-mitigated data-driven circuit learning on noisy quantum hardware. *arXiv:1911.13289 [quant-ph]*, November 2019. arXiv: 1911.13289.
- [32] Peter J. Karalekas, Nikolas A. Tezak, Eric C. Peterson, Colm A. Ryan, Marcus P. da Silva, and Robert S. Smith. A quantum-classical cloud platform optimized for variational hybrid algorithms. *Quantum Science and Technology*, 5(2):024003, April 2020. arXiv: 2001.04449.
- [33] Michael R. Geller and Mingyu Sun. Efficient correction of multiqubit measurement errors. *arXiv:2001.09980 [quant-ph]*, July 2020. arXiv: 2001.09980.
- [34] Michael R Geller. Rigorous measurement error correction. *Quantum Science and Technology*, 5(3):03LT01, June 2020.
- [35] C. Bauer R. Hicks and B. Nachman. Readout Rebalancing for Near Term Quantum Computers. *Phys. Rev. A*, 103:022407, 2021.
- [36] E. F. Dumitrescu, A. J. McCaskey, G. Hagen, G. R. Jansen, T. D. Morris, T. Papenbrock, R. C. Pooser, D. J. Dean, and

- P. Lougovski. Cloud quantum computing of an atomic nucleus. *Phys. Rev. Lett.*, 120:210501, May 2018.
- [37] Suguru Endo, Simon C. Benjamin, and Ying Li. Practical quantum error mitigation for near-future applications. *Phys. Rev. X*, 8:031027, Jul 2018.
- [38] Kristan Temme, Sergey Bravyi, and Jay M. Gambetta. Error mitigation for short-depth quantum circuits. *Phys. Rev. Lett.*, 119:180509, Nov 2017.
- [39] Abhinav Kandala, Kristan Temme, Antonio D. Córcoles, Antonio Mezzacapo, Jerry M. Chow, and Jay M. Gambetta. Error mitigation extends the computational reach of a noisy quantum processor. *Nature*, 567(7749):491–495, 2019.
- [40] Andre He, Benjamin Nachman, Wibe A. de Jong, and Christian W. Bauer. Zero-noise extrapolation for quantum-gate error mitigation with identity insertions. *Phys. Rev. A*, 102:012426, Jul 2020.
- [41] Matthew Otten and Stephen K. Gray. Recovering noise-free quantum observables. *Physical Review A*, 99(1), Jan 2019.
- [42] Seyon Sivarajah, Silas Dilkes, Alexander Cowtan, Will Simmons, Alec Edgington, and Ross Duncan. $t|ket\rangle$: a retargetable compiler for nisq devices. *Quantum Science and Technology*, 6(1):014003, Nov 2020.
- [43] Thomas Häner, Damian S Steiger, Krysta Svore, and Matthias Troyer. A software methodology for compiling quantum programs. *Quantum Science and Technology*, 3(2):020501, Feb 2018.
- [44] Alexander S. Green, Peter LeFanu Lumsdaine, Neil J. Ross, Peter Selinger, and Benoît Valiron. Quipper. *ACM SIGPLAN Notices*, 48(6):333–342, Jun 2013.
- [45] Ali JavadiAbhari, Shruti Patil, Daniel Kudrow, Jeff Heckey, Alexey Lvov, Frederic T. Chong, and Margaret Martonosi. Scaffcc: Scalable compilation and analysis of quantum programs. *Parallel Computing*, 45:2–17, Jun 2015.
- [46] Krysta Svore, Martin Roetteler, Alan Geller, Matthias Troyer, John Azariah, Christopher Granade, Bettina Heim, Vadym Kliuchnikov, Mariia Mykhailova, and Andres Paz. Q#. *Proceedings of the Real World Do-*
- main Specific Languages Workshop 2018 on - RWDSL2018*, 2018.
- [47] Nathan Killoran, Josh Izaac, Nicolás Quesada, Ville Bergholm, Matthew Amy, and Christian Weedbrook. Strawberry fields: A software platform for photonic quantum computing. *Quantum*, 3:129, Mar 2019.
- [48] Gadi Aleksandrowicz et al. Qiskit: An Open-source Framework for Quantum Computing, January 2019.
- [49] Robert S. Smith, Michael J. Curtis, and William J. Zeng. A practical quantum instruction set architecture, 2016.
- [50] Damian S. Steiger, Thomas Häner, and Matthias Troyer. Projectq: an open source software framework for quantum computing. *Quantum*, 2:49, Jan 2018.
- [51] Quantum AI team and collaborators. Cirq, October 2020.
- [52] Alexander J. McCaskey, Eugene F. Dumitrescu, Dmitry Liakh, Mengsu Chen, Wu chun Feng, and Travis S. Humble. A language and hardware independent approach to quantum-classical computing, 2018.
- [53] Prakash Murali, Norbert Matthias Linke, Margaret Martonosi, Ali Javadi Abhari, Nhung Hong Nguyen, and Cinthia Huerta Alderete. Full-stack, real-system quantum computer studies: Architectural comparisons and design insights, 2019.
- [54] Robert S Smith, Eric C Peterson, Erik J Davis, and Mark G Skilbeck. quilc: An optimizing quil compiler, February 2020.
- [55] Yunseong Nam, Neil J. Ross, Yuan Su, Andrew M. Childs, and Dmitri Maslov. Automated optimization of large quantum circuits with continuous parameters. *npj Quantum Information*, 4(1), May 2018.
- [56] Davide Venturelli, Minh Do, Bryan O’Gorman, Jeremy Frank, Eleanor Rieffel, Kyle E. C. Booth, Thanh Nguyen, Parvathi Narayan, and Sasha Nanda. Quantum circuit compilation: An emerging application for automated reasoning. In *Proceedings of the Scheduling and Planning Applications Workshop (SPARK2019)*, 2019.
- [57] Prakash Murali, Jonathan M. Baker, Ali Javadi Abhari, Frederic T. Chong, and Margaret Martonosi. Noise-adaptive com-

- piler mappings for noisy intermediate-scale quantum computers, 2019.
- [58] Prakash Murali, David C. McKay, Margaret Martonosi, and Ali Javadi-Abhari. Software mitigation of crosstalk on noisy intermediate-scale quantum computers. *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, Mar 2020.
- [59] Eric C. Peterson, Gavin E. Crooks, and Robert S. Smith. Two-qubit circuit depth and the monodromy polytope. *Quantum*, 4:247, Mar 2020.
- [60] Nelson Leung, Mohamed Abdelhafez, Jens Koch, and David Schuster. Speedup for quantum optimal control from automatic differentiation based on graphics processing units. *Physical Review A*, 95(4), Apr 2017.
- [61] Pranav Gokhale, Yongshan Ding, Thomas Propson, Christopher Winkler, Nelson Leung, Yunong Shi, David I. Schuster, Henry Hoffmann, and Frederic T. Chong. Partial compilation of variational algorithms for noisy intermediate-scale quantum machines. *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, Oct 2019.
- [62] Nino Shervashidze, Pascal Schweitzer, Erik Jan van Leeuwen, Kurt Mehlhorn, and Karsten M. Borgwardt. Weisfeiler-lehman graph kernels. *J. Mach. Learn. Res.*, 12(null):2539–2561, November 2011.
- [63] Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. Exploring Network Structure, Dynamics, and Function using NetworkX. In Gaël Varoquaux, Travis Vaught, and Jarrod Millman, editors, *Proceedings of the 7th Python in Science Conference*, pages 11 – 15, Pasadena, CA USA, 2008.
- [64] Adriano Barenco, Charles H. Bennett, Richard Cleve, David P. DiVincenzo, Norman Margolus, Peter Shor, Tycho Sleator, John A. Smolin, and Harald Weinfurter. Elementary gates for quantum computation. *Phys. Rev. A*, 52:3457–3467, Nov 1995.
- [65] Bob Coecke and Ross Duncan. Interacting quantum observables: categorical algebra and diagrammatics. *New Journal of Physics*, 13(4):043016, Apr 2011.
- [66] Ross Duncan, Aleks Kissinger, Simon Perdrix, and John van de Wetering. Graph-theoretic simplification of quantum circuits with the zx-calculus. *Quantum*, 4:279, Jun 2020.
- [67] Stephen P. Jordan, Keith S. M. Lee, and John Preskill. Quantum Algorithms for Quantum Field Theories. *Science*, 336:1130–1133, 2012.
- [68] Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA, 2009.
- [69] The following 10 pre-defined passes are considered for the t|ket> optimization: *EulerAngleReduction(OpType.Rz, OpType.Rx)*, *RemoveRedundancies*, *GuidedPauliSimp*, *SquashHQS*, *FlattenRegisters*, *OptimizePhaseGadgets*, *KAKDecomposition*, *USquashIBM*, *CliffordSimp*, *FullPeepholeOptimize*. Two more passes, *RebaseIBM*, *CommuteThroughMultis*, are also used once before selecting the pass from the list.
- [70] IBM Quantum, 2021.
- [71] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [72] Kathleen E. Hamilton, Tyler Kharazi, Titus Morris, Alexander J. McCaskey, Ryan S. Bennink, and Raphael C. Pooser. Scalable quantum processor noise characterization, 2020.