

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

Abstract. 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 quantum devices produced over a next decade. 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 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. 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.

One of the general strategies for improving the performance of NISQ computers 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` [16], 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, where CNOT is a controlled Pauli-X gate. There are also a variety of other toolkits for circuit optimization, including hardware-specific packages for quantum circuits (see e.g., [17–20]). 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. The second technique optimizes a generic circuit by eliminating unnecessary gates or unused qubits such that the circuit depth¹ 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 (Toffoli gate is a controlled-controlled Pauli-X gate). 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, (pronounced “excel”) standing 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 useful benchmark because it provides an exponential speedup over the most efficient know 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.

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. 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.

¹The 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.




Figure 1. Possible RSG patterns for a given target RSG corresponding to the three levels of matching criteria.

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. Such a structure can be described by using a directed acyclic graph (DAG). The 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 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). The gate set pattern recognition can then 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 using depth-first search with heuristic pruning, and (2) checking the DAG isomorphism by graph hashing with Weisfeiler Lehman graph hash [21], as implemented in the NetworkX library [22].

2.1.2 Tiered extraction of recurring gate sets

The appearance pattern of RSGs in a 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 gate recurrence:

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, that is, the RSGs must appear on an identical set of qubits in the circuit. The appearance patterns of the RSGs are illustrated in Fig. 1 for the three matching criteria. The Level 2 is used in the

following studies, but the different levels will provide further opportunities for gate reduction beyond what has been done so far.

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 specified number of top-ranked RSGs are extracted from the circuit in this step.

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 shallow 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 would provide a shorter, more efficient circuit. Compared to Level 1, the Level 2 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 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 $|j\rangle_{\text{ctl}}$ 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, corresponding to the state $|j\rangle_{\text{ctl}} = |11 \dots 1\rangle = |2^m - 1\rangle_{\text{ctl}}$. This allows one to classify the state of the system into three general classes with 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 a certain condition.

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 $O(m)$ Toffoli and controlled- U gates [23]. 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 $\tilde{N} = nN$ 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 $O(\tilde{N})$ 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 $O(2^n)$ 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 $O(M)$ operations. Repeating this for all \tilde{N} gates requires $O(\tilde{N}2^n)$ for the classical simulation and $O(\tilde{N}^2 M)$ when using quantum measurements.

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.

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 can be determined with $O(\tilde{N}^2 M)$ operations, one can compute the probabilities for each possible bitstring, and therefore decide whether to simplify a controlled operation using $O(\tilde{N})$ operations.

Note that superfluous controlled operations can also be found and eliminated using the so-called ZX calculus [24, 25]. 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, 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 below.

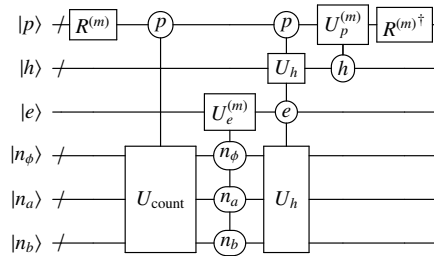


Figure 2. The m^{th} step of the quantum circuit for the algorithm proposed in Ref. [15]. There are three physical registers: $|p\rangle$ containing the set of particles at this step; $|h\rangle$ for the branching history; and $|e\rangle$ which is a binary variable representing the presence or absence 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 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.

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 [26]. 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 $f_1\bar{f}_2$ (\bar{f}_1f_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 2 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 $f_a\bar{f}_b$ (\bar{f}_af_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 accordingly. 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.

3.2 Experimental setup

The QPS simulation is implemented into a quantum circuit using IBM Qiskit version 0.22.0 [27] with Terra 0.15.2, Aer 0.6.1 and Ignis 0.4.0 APIs in Python 3.8 [28]. 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>` [16] 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 passes². The passes are tried one 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|ket>` performance in the remainder of the studies.

The QPS algorithm is executed on the 27-qubit “`ibmq_sydney`” [29], 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 constraints³. 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 considered⁴.

First, the RSG pattern recognition is performed against the circuit. When the Level 2 RSG pattern recognition is applied, two RSGs are identified 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. 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),

²The following 10 pre-defined passes are considered for the `t|ket>` optimization: *EulerAngleReduction(OpType.Rz,OpType.Rx)*, *RemoveRedundancies*, *GuidedPauliSimp*, *SquashHQIS*, *FlattenRegisters*, *OptimizePhaseGadgets*, *KAKDecomposition*, *USquashIBM*, *CliffordSimp*, *FullPeepholeOptimise*. Two more passes, *RebaseIBM*, *CommuteThroughMultis*, are also used once before selecting the pass from the list.

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

⁴Ref. [15] noted that when these are unphysically removed, the circuit can be simulated efficiently classically (see also Ref. [14])




Figure 3. 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.

and removing unused qubits. Nonzero-amplitude computational basis states are identified through classical calculation.

After the optimization, the quantum gates in the circuit are decomposed into single-qubit gates (U_1, U_2, U_3) and CNOT gates. Figure 3 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 figure compares the values from the original circuit and the circuits optimized with t|ket) only, AQCEL only and 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), 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) in reducing the gate counts, and 46% more specifically for CNOT, and makes the circuit 45% shorter. Combining the two optimizers (AQCEL applied first, then t|ket)), 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) alone for gate reduction (49% for CNOT only), producing a 50% shorter circuit.

For the AQCEL optimizer, the gate reduction occurs mostly at the stage where the redundant qubit controls are removed. Starting with 1279 gates (excluding barrier and measurement gates), the first adjacent gate-pair elimination, the redundant qubit control reduction, and the second gate-pair elimination steps remove 170, 510 (40% of the 1279 gates), and 6 gates, respectively. 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 sub-dominant contribution of 7% from the redundant qubit control reduction. The number of qubits finally reduces from 24 to 21 with AQCEL, while it does not change with t|ket).

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

3.3.2 Circuit optimization for $N_{evol} = 1$ 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. There-




Figure 4. 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.

fore, we consider the circuit with only one branching step using the `sydney` and the `statevec` 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, consistent with the exponential behavior of the computing cost of the step. The number of qubits is reduced from 15 to 13 with AQCEL. 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 4 shows the same set of distributions as in Fig. 3, but for the one-branching step QPS simulation with the `sydney`-specific transpilation. The AQCEL optimizer reduces the number of native gates significantly in this case, too. 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 optimizer performance 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 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 gate errors accumulate throughout the circuit execution and are difficult to correct. Instead, in AQCEL, we opt to ignore the observed bitstrings with occurrence below certain “cut-off” thresholds, under the assumption that the gate errors act as a perturbation that inserts spurious computational basis states with small amplitudes into the system. This can be improved in future with additional computational




Figure 5. 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.

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 the $U_{1,2,3}$ and 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:

$$p_\epsilon = 1 - \left[\prod_i (1 - \epsilon_U^{(i)})^{n_U^{(i)}} \prod_{i \neq j} (1 - \epsilon_{CX}^{(i,j)})^{n_{CX}^{(i,j)}} \right] \sim N_{CX} \epsilon_{CX}, \quad (2)$$

where $n_U^{(i)}$ and $n_{CX}^{(i,j)}$ are the number of $U_{1,2,3}$ and CNOT gates acting on the corresponding qubits, respectively. In the last approximation, we have assumed that all CNOT errors are equal, much larger than single gate errors but still much much smaller than one: $\epsilon_U^{(i)} \ll \epsilon_{CX}^{(i,j)} = \epsilon_{CX} \ll 1$. The first cutoff threshold is $s_\epsilon^{\text{high}} := p_\epsilon$, 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, $s_\epsilon^{\text{low}} := p_\epsilon / 2^m$, 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, $s_\epsilon^{\text{med}} := (s_\epsilon^{\text{low}} + s_\epsilon^{\text{high}}) / 2$.

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).

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 results in an elimination of most of the controlled gates in the later part of the circuit, rendering the circuit practically mean-




Figure 6. 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} .

ingless. Therefore, the actual threshold to be used with AQCEL should be chosen considering the tradeoff between the circuit optimization efficiency and the accuracy of the optimized circuit⁵.

Figure 5 shows the gate counts obtained from AQCEL optimizations using actual measurements on the `sydney` under the dynamic cutoff thresholds. The gate counts decrease as the threshold is raised from $s_{\epsilon}^{\text{low}}$ to $s_{\epsilon}^{\text{high}}$, as expected. The same distributions obtained with the static thresholds (not shown) indicate that almost no gate survives under the threshold of 40%, likely implying a significant loss of accuracy of the computation result. The number of qubits is reduced from 15 to 13 with the threshold of $s_{\epsilon}^{\text{low}}$ or $s_{\epsilon}^{\text{med}}$, and to 11 with $s_{\epsilon}^{\text{high}}$. Under the 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 $s_{\epsilon}^f \geq 35\%$.

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

$$F = \sum_k \sqrt{p_k^{\text{orig}} p_k^{\text{opt}}}, \quad (3)$$

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 unity of F_{sim} 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 quantum computer measurements. The p^{opt} is estimated from the rate of a bitstring occurring in a large number of repeated measurements. The p^{orig} is computed using simulation as for the F_{sim} . Even if the optimized

⁵In 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 7. 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} .

circuit is identical to the original circuit, the presence of noise will make $F_{\text{meas}} < 1$, with the discrepancy from unity 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 and raise the F_{meas} value. However, in some cases the CNOT gate removal affects low-amplitude computational basis states, meaning that the optimized circuit can differ from the original circuit, hence suppress the F_{meas} value. Thus, the F_{meas} reflects the tradeoff of making the circuit shorter and changing the circuit through optimization.

Figure 6 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 does increase 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.

When the elimination of redundant qubit controls is performed based on measurements using quantum computer, the F_{meas} versus gate counts become those shown in Fig. 7 with the dynamic thresholds of $s_{\epsilon}^{\text{high}}$, $s_{\epsilon}^{\text{med}}$, and $s_{\epsilon}^{\text{low}}$. The results with the s_{ϵ}^f thresholds (not shown) show that the F_{meas} first increases but then gets worse, signaling that the optimized circuit becomes too far from the original circuit. For the circuit considered here, the optimization performance seems best for $s_{\epsilon}^f \sim 25\%$, therefore the capped variants are also shown in Fig. 7 with the 25% threshold. 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 different choices of cutoff thresholds are summarized in Figs. 8 and 9. It is worth noting that most of the AQCEL-based optimization 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 result but the existence of the controlled gates produces those spurious states




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 under different schemes.




Figure 9. 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} .

under the effect of hardware noise. An exact identification of computational basis states with nonzero amplitudes using classical simulation does not allow removing such qubit controls, effectively degrading the F_{meas} .

Figure 10 shows the 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 better than the original circuit because gate reductions suppress the effect of hardware noise.

4 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

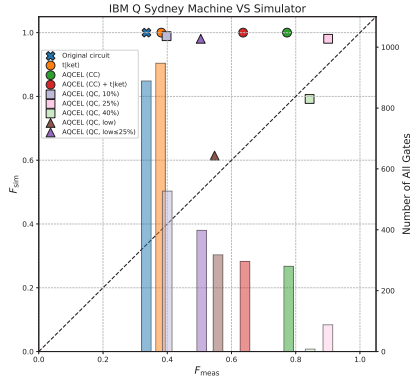


Figure 10. Fidelities F_{meas} versus F_{sim} for the one-branching step QPS circuit transpiled considering the `ibmq_sydney` topology before and after optimization under different schemes. These transpiled circuits are executed on the `sydney` to obtain the F_{meas} and a statevector simulator to obtain the F_{sim} . The vertical bars represent the total number of gates in the circuits (gauged by the right-hand vertical axis).

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 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 `tket`, which is one of the industry-standard optimization tools, while retaining the accuracy of the probability distributions of the final state.

The AQCEL protocol 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. [30] and related approaches that use fewer resources [31–35] or further mitigate the errors [36]. 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. 9, 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.

References

- [1] J. Preskill, *Quantum* **2**, 79 (2018)
- [2] A. Mott, J. Job, J.R. Vlimant, D. Lidar, M. Spiropulu, *Nature* **550**, 375 (2017)
- [3] A. Zlokapa, A. Mott, J. Job, J.R. Vlimant, D. Lidar, M. Spiropulu (2019), [1908.04480](https://arxiv.org/abs/1908.04480)
- [4] J. Chan, W. Guan, S. Sun, A.Z. Wang, S.L. Wu, C. Zhou, M. Livny, F. Carminati, A.D. Meglio, *PoS LeptonPhoton2019*, 049 (2019)
- [5] K. Terashi, M. Kaneda, T. Kishimoto, M. Saito, R. Sawada, J. Tanaka, *Comput. Softw. Big Sci.* **5**, 2 (2020)
- [6] W. Guan, G. Perdue, A. Pesah, M. Schuld, K. Terashi, S. Vallecorsa, J.R. Vlimant, *Machine Learning: Science and Technology* (2020), [2005.08582](https://arxiv.org/abs/2005.08582)
- [7] A. Zlokapa, A. Anand, J.R. Vlimant, J.M. Duarte, J. Job, D. Lidar, M. Spiropulu (2019), [1908.04475](https://arxiv.org/abs/1908.04475)
- [8] C. Tüysüz, F. Carminati, B. Demirköz, D. Dobos, F. Fracas, K. Novotny, K. Potamianos, S. Vallecorsa, J.R. Vlimant, *EPJ Web Conf.* **245**, 09013 (2020)
- [9] I. Shapoval, P. Calafiura, *EPJ Web Conf.* **214**, 01012 (2019)
- [10] F. Bapst, W. Bhimji, P. Calafiura, H. Gray, W. Lavrijsen, L. Linder, *Comput. Softw. Big Sci.* **4**, 1 (2019)
- [11] A.Y. Wei, P. Naik, A.W. Harrow, J. Thaler, *Phys. Rev. D* **101**, 094015 (2020)
- [12] S. Das, A.J. Wildridge, S.B. Vaidya, A. Jung (2019), [1903.08879](https://arxiv.org/abs/1903.08879)
- [13] K. Cormier, R. Di Sipio, P. Wittek, *JHEP* **11**, 128 (2019)
- [14] D. Provasoli, B. Nachman, C. Bauer, W.A. de Jong, *Quantum Science and Technology* **5**, 035004 (2020)
- [15] C.W. Bauer, B. Nachman, D. Provasoli, W.A. De Jong, *Phys. Rev. Lett.* **126**, 062001 (2021), [1904.03196](https://arxiv.org/abs/1904.03196)
- [16] S. Sivarajah, S. Dilkes, A. Cowtan, W. Simmons, A. Edgington, R. Duncan, *Quantum Science and Technology* **6**, 014003 (2020)
- [17] T. Häner, D.S. Steiger, K. Svore, M. Troyer, *Quantum Science and Technology* **3**, 020501 (2018)
- [18] Quantum AI team and collaborators, *Cirq* (2020), <https://doi.org/10.5281/zenodo.4062499>
- [19] R.S. Smith, E.C. Peterson, E.J. Davis, M.G. Skilbeck, *quilc: An Optimizing Quil Compiler* (2020), <https://doi.org/10.5281/zenodo.3677537>
- [20] P. Murali, J.M. Baker, A.J. Abhari, F.T. Chong, M. Martonosi, *Noise-adaptive compiler mappings for noisy intermediate-scale quantum computers* (2019), [1901.11054](https://arxiv.org/abs/1901.11054)
- [21] N. Shervashidze, P. Schweitzer, E.J. van Leeuwen, K. Mehlhorn, K.M. Borgwardt, *J. Mach. Learn. Res.* **12**, 2539–2561 (2011)
- [22] A.A. Hagberg, D.A. Schult, P.J. Swart, *Exploring Network Structure, Dynamics, and Function using NetworkX*, in *Proceedings of the 7th Python in Science Conference*, edited by G. Varoquaux, T. Vaught, J. Millman (Pasadena, CA USA, 2008), pp. 11 – 15
- [23] A. Barenco, C.H. Bennett, R. Cleve, D.P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J.A. Smolin, H. Weinfurter, *Phys. Rev. A* **52**, 3457 (1995)

- [24] B. Coecke, R. Duncan, *New Journal of Physics* **13**, 043016 (2011)
- [25] R. Duncan, A. Kissinger, S. Perdrix, J. van de Wetering, *Quantum* **4**, 279 (2020)
- [26] S.P. Jordan, K.S.M. Lee, J. Preskill, *Science* **336**, 1130 (2012)
- [27] G. Aleksandrowicz et al., *Qiskit: An Open-source Framework for Quantum Computing* (2019), <https://doi.org/10.5281/zenodo.2562111>
- [28] G. Van Rossum, F.L. Drake, *Python 3 Reference Manual* (CreateSpace, Scotts Valley, CA, 2009), ISBN 1441412697
- [29] *IBM Quantum*, <https://quantum-computing.ibm.com> (2021), accessed: 2020-02-27
- [30] C.W. Bauer, W.A. De Jong, B. Nachman, M. Urbanek, *npj Quantum Inf.* **6**, 84 (2020)
- [31] M.R. Geller, M. Sun (2020), [2001.09980](https://arxiv.org/abs/2001.09980)
- [32] C. Song, K. Xu, W. Liu, C.p. Yang, S.B. Zheng, H. Deng, Q. Xie, K. Huang, Q. Guo, L. Zhang et al., *Physical Review Letters* **119**, 180511 (2017)
- [33] M. Gong, M.C. Chen, Y. Zheng, S. Wang, C. Zha, H. Deng, Z. Yan, H. Rong, Y. Wu, S. Li et al., *Physical Review Letters* **122**, 110501 (2019)
- [34] K.X. Wei, I. Lauer, S. Srinivasan, N. Sundaresan, D.T. McClure, D. Toyli, D.C. McKay, J.M. Gambetta, S. Sheldon, *Physical Review A* **101**, 032343 (2020)
- [35] K.E. Hamilton, T. Kharazi, T. Morris, A.J. McCaskey, R.S. Bennink, R.C. Pooser (2020), [2006.01805](https://arxiv.org/abs/2006.01805)
- [36] R. Hicks, C.W. Bauer, B. Nachman, *Phys. Rev. A* **103**, 022407 (2021)