

Initial-State Dependent Optimization of Controlled Gate Operations with Quantum Computer

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

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, connectivity, and coherence times, a quantum circuit optimization is essential to make the best use of near-term quantum devices. We introduce a new circuit optimizer called AQCEL, which aims to remove redundant controlled operations from controlled gates, depending on initial states of the circuit. Especially, the AQCEL can remove unnecessary qubit controls from multi-controlled gates in polynomial computational resources, even when all the relevant qubits are entangled, by identifying zero-amplitude computational basis states using a quantum computer. As a benchmark, the AQCEL is deployed on a quantum algorithm designed to model final state radiation in high energy physics. For this benchmark, we have demonstrated that the AQCEL-optimized circuit can produce equivalent final states with much smaller number of gates. Moreover, when deploying AQCEL with a noisy intermediate scale quantum computer, it efficiently produces a quantum circuit that approximates the original circuit with high fidelity by truncating low-amplitude computational basis states below certain thresholds. Our technique is useful for a wide variety of quantum algorithms, opening up new possibilities to further simplify quantum circuits to be more effective for real devices.

Wonho Jang: jang@icepp.s.u-tokyo.ac.jp

Koji Terashi: koji.terashi@cern.ch

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, qubit connectivity and capability to stay coherent for the entirety of 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, 3, 4, 5, 6, 7], reconstructions of charged particle trajectories [8, 9, 10, 11] and physics objects [12, 13], unfolding measured distributions [14] as well as simulation of multi-particle emission processes [15, 16]. 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 [17, 18, 19, 20, 21, 22] and gate errors can be mitigated by systematically enlarging errors before extrapolating to zero error [23, 24, 25, 26, 27, 28]. A complementary strategy to error mitigation is quantum compilation. 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. Widely used tools are Qiskit [29] and t|ket) [30], which contain a variety of architecture-agnostic and architecture-specific routines. There are also a variety of other toolkits for circuit optimization, including hardware-specific packages for quantum circuits [31, 32, 33, 34, 35, 29, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49].

Among the gates used for an algorithm encoding, multi-controlled gates are significant error sources because they result in many CNOT gates after the decomposition [50], and also require SWAP gates to fit within limited qubit topology. The costs to implement multi-controlled gates can be reduced by using relative phase Toffli gates [51], ancilla qubits [52] or qutrits [53, 54, 55, 56, 57, 58, 59] in the implementation.

An alternative approach for reducing the costs is to remove unnecessary qubit controls. The reversible circuit synthesis can reduce redundant qubit controls while maintaining the equivalence of a quantum circuit before and after the optimization [60]. Previous work with reversible circuit synthesis has largely focused on circuits composed of CNOT gates [61]. The circuit synthesis is extended later to more general quantum circuits, e.g., those composed of CNOTs and Z -basis rotation gates [62]. It is possible to remove, beyond reversible circuit synthesis, more unnecessary controlled operations if we consider maintaining the equivalence of the final state. Imagine that there is a n -qubit quantum circuit designed to work with different initial states and it is executed with a given initial state such as $|0\rangle^{\otimes n}$. In this case, the circuit will reach only a selected set of intermediate states and some operations may become trivial. Such initial-state dependent circuit optimization may find more rooms for optimization if the equivalence of the final state, not the circuit itself, is preserved. Thus, it will enable more aggressive reduction of unnecessary controlled operations than initial-state independent circuit optimization.

There are two main approaches for initial-state dependent circuit optimization. The first one is a continuous optimization that trains an ansatz with parameters [63, 64]. The second one is a discretized optimization in which some controlled operations are removed from a gate if the quantum state satisfies a specific condition at the point

where the gate is operated [49, 65, 66]. We focus on the latter in this paper. Among existing discretized optimization protocols that account for initial states, the Relaxed Peephole Optimization (RPO) [49] reduces controlled operations when the qubits in the X - or Z -basis states are used as control qubits or a target qubit of a controlled gate. This protocol, however, cannot remove qubit controls in the case where all relevant qubits are entangled. The ZX-calculus [65, 66] exploits a copy-rule for removing a qubit control from CNOT gate when the control qubit state is $|1\rangle$. This leads to an initial-state dependent circuit optimization, but it cannot remove all qubit controls within polynomial complexity ¹.

The novel optimization protocol proposed in this paper has three distinct features. First, there is the ability of removing redundant qubit controls no matter whether all the relevant qubits are entangled or not. Second, the identification of zero- or low-amplitude computational basis states using a quantum computer allows one to obtain bitstrings in polynomial time, otherwise exponential resources are required in the classical calculations. Third, the decomposition of multi-controlled U gates into Toffoli gates and singly-controlled U gates enables us to perform the search of all unnecessary qubit controls in polynomial resources. This new optimization protocol also serves a new efficient method to approximate quantum circuits in the NISQ era by truncating low-amplitude computational basis states that do not contribute significantly to the final state.

This optimization protocol is called AQCEL (and pronounced “excel”) for *Advancing Quantum Circuit* by ICEPP and LBNL. To demonstrate the effectiveness of the AQCEL protocol, we will use a quantum algorithm that models a *parton shower* [16]. This algorithm provides a useful benchmark because it is designed to work with different initial states corresponding to different initial particles, meaning that the quantum circuits have redundancy for a specific initial state.

This paper is organized as follows. Section 2 provides an overview of the AQCEL protocol. The

¹In fact, the ZX-calculus is complete in the formal logic sense of the word, such that one can always prove that all unnecessary qubit controls can be removed using rules of the ZX-calculus [67]. However, in general this scheme requires exponential resources. Nevertheless, the ZX-calculus is still incredibly powerful and underlies many of the optimization techniques of quantum transpilers.

application of this protocol to the HEP example is presented in Sec. 3. Following a brief discussion about the applicability and future extensions of the protocol in Sec. 4, the paper concludes in Sec. 5.

2 AQCEL optimization protocol

First, we summarize the concept of removing redundant controlled operations from controlled gates depending on initial states of a quantum circuit. Then, the AQCEL protocol of removing redundant qubit controls and the methods for executing the whole optimization in polynomial resources are described.

2.1 Basic idea of redundant controlled operations removal

A controlled gate performs a different operation depending on the quantum state at the point where the gate is applied. Let m be the number of control qubits of this gate. 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 unlabeled ket corresponds to the rest of the system. We write the states as integers with $0 \leq j \leq 2^m - 1$ and $0 \leq k \leq 2^{n-m} - 1$. We assume that the controlled gate is applied to computational basis states whose bitstrings on all control qubits are 1, which corresponds 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 using the amplitudes $c_{j,k}$:

Triggering : $c_{j,k} \neq 0$ if and only if $j = 2^m - 1$.

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, and that is our interest.

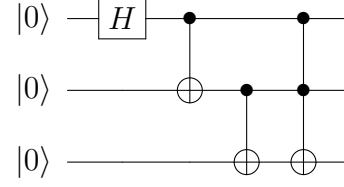


Figure 1: A quantum circuit in $|0\rangle^{\otimes 3}$ initial state. The control and target qubits for the Toffoli gate are entangled because the Hadamard and CNOT gates create the GHZ state.

As an example of this concept, consider the simple circuit in Fig. 1 composed of three qubits in $|0\rangle^{\otimes 3}$ initial state. At the Toffoli gate, the quantum state is in the superposition of $|000\rangle$ and $|111\rangle$, which is a Greenberger-Horne-Zeilinger state where the qubits are maximally entangled. This is an undetermined state for the Toffoli gate. Moreover, all the control qubits and the target qubit are entangled, which is a difficult case for the qubit control reduction. However, the AQCEL can remove one of the two qubit controls from the Toffoli, replacing it with a CNOT gate controlled only by the remaining one.

2.2 General conditions to eliminate qubit controls

Given a multi-qubit controlled- U gate with m control qubits, denoted by $C^m[U]$, and a system in an undetermined state $|\psi\rangle$ defined in Sec. 2.1, we can derive general conditions for a part of the controlled operations to be removed, as follows.

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

$$|\psi\rangle = \sum_{i,l,k} \tilde{c}_{i,l,k} |i\rangle_{\text{ctl}'} \otimes |l\rangle_{\text{free}} \otimes |k\rangle, \quad (2)$$

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),

$$|i\rangle_{\text{ctl}'} \otimes |l\rangle_{\text{free}} = |2^x i + l\rangle_{\text{ctl}}, \quad (3)$$

and therefore

$$\tilde{c}_{i,l,k} = c_{2^x i + l, k}. \quad (4)$$

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

$$C^m[U]|\psi\rangle = \sum_{j=0}^{2^m-2} \sum_k c_{j,k} |j\rangle |k\rangle + \sum_k c_{2^m-1,k} |2^m-1\rangle U|k\rangle, \quad (5)$$

where ket subscripts and the tensor product symbols are omitted for simplicity. In contrast, a new gate with fewer controls gives

$$C^{m-x}[U]|\psi\rangle = \sum_{i=0}^{2^{m-x}-2} \sum_{l,k} \tilde{c}_{i,l,k} |i\rangle |l\rangle |k\rangle + \sum_{l,k} \tilde{c}_{2^{m-x}-1,l,k} |2^{m-x}-1\rangle |l\rangle U|k\rangle. \quad (6)$$

For the removal of x qubit controls to be allowed, the right hand sides of Eqs. (5) and (6) must be identical. This requires

$$\sum_{l=0}^{2^x-2} \sum_k \tilde{c}_{2^{m-x}-1,l,k} |2^{m-x}-1\rangle |l\rangle U|k\rangle = \sum_{l=0}^{2^x-2} \sum_k c_{2^m-2^x+l,k} |2^m-2^x+l\rangle |k\rangle. \quad (7)$$

Denoting

$$U|k\rangle = \sum_{k'} u_{kk'} |k'\rangle \quad (8)$$

and recalling Eq. (3), Eq. (4), Eq. (7) implies (replacing $k' \leftrightarrow k$ on the left hand side)

$$\sum_{l=0}^{2^x-2} \sum_{k,k'} \tilde{c}_{2^{m-x}-1,l,k'} u_{k'k} |2^{m-x}-1\rangle |l\rangle |k\rangle = \sum_{l=0}^{2^x-2} \sum_k \tilde{c}_{2^{m-x}-1,l,k} |2^{m-x}-1\rangle |l\rangle |k\rangle. \quad (9)$$

Then, we have

$$\sum_{k'} \tilde{c}_{2^{m-x}-1,l,k'} u_{k'k} = \tilde{c}_{2^{m-x}-1,l,k} \quad \forall l \in \{0, 1, \dots, 2^x-2\}, k. \quad (10)$$

Eq. (10) 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 $0 \leq l \leq 2^x-2$, or if $\tilde{c}_{2^{m-x}-1,l,k} = 0$ for $0 \leq l \leq 2^x-2$ and all k .

Since the cost of exactly computing the complex amplitudes of the quantum state is exponential, in AQCEL we only consider this second condition:

$$\tilde{c}_{2^{m-x}-1,l,k} = 0 \quad \forall l \in \{0, 1, \dots, 2^x-2\}, k. \quad (11)$$

This removal of redundant qubit controls therefore requires us to find out if $|\psi\rangle$ satisfies Eq. (11) for each controlled gate.

In the previous optimization based on RPO, if the target qubit is assumed to be entangled, the condition for the removal of qubit controls from $C^m[U]$ is that control qubits which will be removed must be $|1\rangle$ in a pure state. The condition can be written as

$$\tilde{c}_{i,l,k} = 0 \quad \forall l \in \{0, 1, \dots, 2^x-2\}, i, k. \quad (12)$$

Eq. (12) is a sufficient condition of Eq. (11), which means that the AQCEL has more rooms for removal of unnecessary qubit controls. A difference between Eq. (11) and Eq. (12) resides in whether one can remove unnecessary qubit controls from entangled control qubits or not, under the condition that target qubit is also entangled.

2.3 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 nonzero-amplitude basis states depends on the initial state. This is why the three classes of the states on control qubits arise.

As mentioned in Sec. 2.2, we have to find out if $|\psi\rangle$ satisfies Eq. (11). It can be written down as

$$\sum_k |\tilde{c}_{2^{m-x}-1,l,k}|^2 = 0 \quad \forall l \in \{0, 1, \dots, 2^x-2\}. \quad (13)$$

Eq. (13) requires that there is no computational basis state whose bitstring on all control qubits of $C^{m-x}[U]$ is $|11 \dots 1\rangle$, except when the bitstring on the removed x control qubits is also $|11 \dots 1\rangle$. In other words, there should be no bitstring by which $C^m[U]$ is not triggered but $C^{m-x}[U]$ is. This can be verified by the identification of bitstrings on all control qubits of $C^m[U]$.

The possible bitstrings on control qubits at each controlled gate can be determined either through a classical simulation or by measuring the control qubits by a quantum computer repeatedly. In the case of a classical simulation, one performs the full calculation of the amplitudes. 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 ².

A few notes should be taken on the computational costs of the two methods. Consider an n -qubit circuit with N 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 m control qubits M times on each controlled gates by a quantum computer only requires $\mathcal{O}(MN^2 + mMN)$ operations which scales only polynomially with the number of qubits. More details on the estimates of the computational resource necessary for the identification of computational basis states are described in Appendix B.

Note that for noisy quantum computers the measurements of the bitstrings will not be exact due to hardware noise. The list of observed bitstrings would contain contributions from errors on the preceding gates and the measurement itself. In AQCEL, we obtain the calibration matrix for the control qubits (with 8192 shots per measurement) using Qiskit Ignis API [29]. The matrix is then applied to the observed distribution with a least-squares fitting approach. To deal with remaining error contributions after the measurement error mitigation, we opt to ignore the observed bitstrings with occurrence below certain thresholds ³. 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 ⁴.

In order to choose the thresholds, we consider

²When the number of measurements is not enough for a part of low-amplitude computational basis states, the AQCEL approximates the original circuit to a simpler circuit, but this usually loses the identity of the final state. There is a similar effect when low-amplitude computational basis states below thresholds are truncated, as described later.

³In the actual implementation, the threshold of 0.005 is always applied to suppress contributions from imperfect measurement error mitigation.

⁴This is justified under the assumption that the residual gate errors act as a perturbation, inserting spurious computational basis states with small amplitudes into the superposition of the system.

gate errors in the single-qubit gates and CNOT gates ⁵. Let the single-qubit gate 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 approximate the probabilities, p_u and p_{cx} , of measuring the bitstrings without any single-qubit gate or CNOT gate errors occurring anywhere in the circuit by performing qubit-wise (index-dependent) multiplications of the error rates:

$$p_u = \prod_i \left(1 - \epsilon_u^{(i)}\right)^{n_u^{(i)}}, \quad (14)$$

$$p_{cx} = \prod_{i \neq j} \left(1 - \epsilon_{cx}^{(i,j)}\right)^{n_{cx}^{(i,j)}}, \quad (15)$$

where $n_u^{(i)}$ and $n_{cx}^{(i,j)}$ are the numbers of single-qubit gates and CNOT gates acting on the corresponding qubits, respectively. The probability p_ϵ of measuring the bitstrings with at least one gate error occurring anywhere in the circuit is

$$p_\epsilon = 1 - p_u p_{cx} \\ \sim N_{cx} \epsilon_{cx}. \quad (16)$$

In the last approximation, we have assumed that all CNOT errors are equal, and much larger than single-qubit gate errors but still much smaller than one: $\epsilon_u^{(i)} \ll \epsilon_{cx}^{(i,j)} = \epsilon_{cx} \ll 1$. From the p_ϵ , the first threshold is chosen to be

$$s_\epsilon^{\text{dyn}} := p_\epsilon / 2^m, \quad (17)$$

where m is the number of the measured control qubits. This choice of dynamical threshold is motivated by assuming that the quantum errors would result in a uniform distribution of all possible bitstrings according to a depolarizing error model. It should be noted that the p_ϵ increases as the circuit execution proceeds because the p_ϵ accounts for the error rates from all the preceding gates in the circuit. As an alternative strategy to the dynamical threshold, we also examine the static thresholds, s_ϵ^f , that are kept constant throughout the circuit, with the values between 0.005 and 0.3.

Discarding all bitstrings with occurrence under certain thresholds usually modifies the final state of the optimized circuit from one of the original circuit. On the other hand, applying certain

⁵The reported error rates at the time of the experiment, measured during the preceding calibration run of the hardware, are used for the threshold calculation.

thresholds will leave high-amplitude computational basis states, while rejecting low-amplitude computational basis states which do not contribute to the final result meaningfully⁶. Thus, with the proper thresholds, one can produce a quantum circuit which well approximates the original circuit by removing unimportant qubit controls that trigger on low-amplitude computational basis states. In other words, the actual threshold of AQCEL should be selected by considering the trade-off between the noise resilience and the identity of the final state to the original ideal state.

2.4 Elimination of redundant controlled operations

Once the nonzero-amplitude computational basis states are identified at each controlled gate, the next step is to figure out which qubit controls can be removed using Eq. (11). The computational cost of determining the removal of redundant qubit controls would be at most $\mathcal{O}(Mm4^mN)$ (Appendix B), which scales exponentially with m , the number of control qubits (N is the number of multi-qubit controlled gates in the circuit). To avoid this, a generic multi-qubit controlled gate should be decomposed into controlled gates with small, fixed number of control qubits. An arbitrary multi-qubit controlled- U gate with m control qubits can be decomposed into $\mathcal{O}(m)$ Toffoli and controlled- U gates [50]. Besides, these Toffoli gates can be replaced with relative phase implementation of Toffoli gates (referred to as just ‘‘Toffoli’’ hereafter) [51], which reduces the CNOT counts from 6 (in the regular Toffoli decomposition) to 3. Therefore, in the AQCEL scheme, we assume that all controlled gates in a quantum circuit are reduced to Toffoli gates denoted as $C^2[X]$ and singly-controlled unitary operations denoted as $C[U]$. This results in a significant reduction of computational cost of the decision of all redundant qubit controls, because all controlled gates have either 1 or 2 control qubits. However, when controlled gates are decomposed, AQCEL would lose a part of opportunity to remove redundant qubit controls. More details about the decom-

position are described in Appendix A. Since a multi-qubit controlled gate is decomposed into a set of Toffoli gates and its mirror for uncomputation except the central gate, the optimization of controlled operations for the set of Toffoli gates can be applied to the uncomputation part as well.

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 nN controlled gates. With nN gates, the cost for identifying computational basis states (Sec. 2.3) increases up to $\mathcal{O}(n^2N^2M)$ when measuring with a quantum computer. However, the cost for removing unnecessary qubit controls improves from the above exponential scaling to $\mathcal{O}(MnN)$. More details about the resource scaling are given in Appendix B.

After the decomposition, a $C[U]$ gate can be a single unitary U gate if the probability of observing $|1\rangle$ of the control qubit is 1, or removed if the probability is 0. In all other cases, the $C[U]$ gate is kept. For a $C^2[X]$ gate, the similar control reduction can be performed with the probabilities of the four possible states $|00\rangle$, $|01\rangle$, $|10\rangle$ and $|11\rangle$. If the probability of the state $|01\rangle$ ($|10\rangle$) is zero, one can eliminate the first (second) control from the $C^2[X]$ gate (see Eq. (11)). 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 results then
      if  $\{0\}$  is not observed in the measurement 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-

```

⁶Here the low-amplitude computational basis states mean a part of the original computational basis states composing of the ideal quantum state of interest, not those from quantum noise.

sults **then**
if neither $\{00\}$, $\{01\}$ nor $\{10\}$ is observed in the measurement results **then**
turn g into an X gate acting on the target qubit
else if $\{01\}$ is not observed in the measurement results **then**
eliminate the control on q_1
else if $\{10\}$ is not observed in the measurement results **then**
eliminate the control on q_2
end if
else
eliminate g
end if
end if
end for

3 Application to quantum algorithm

The AQCEL optimization protocol described in Sec. 2 has been deployed to the quantum parton shower (QPS) algorithm [16]. We show experimental results from a simulator and quantum hardware, and discuss the optimization performance in terms of the number of CNOT gates and the identity of the final states. The main purpose of this section is to demonstrate that AQCEL optimization works with the determination of bitstrings by a noisy intermediate scale quantum computer in polynomial resources, not to compare with other initial-state dependent optimization protocols.

3.1 Quantum parton shower algorithm

The QPS algorithm in Ref. [16] can start with a fermion that is either type f_1 or f_2 . These fermions can radiate a scalar particle ϕ or not at a given showering step. 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, where the antifermion is denoted by \bar{f} . The shower evolution process is simulated by repeating the step by N_{evol} times. When N_{evol} is small, only a small number of particles are simulated, hence the number of non-zero computational basis states is also small. In addition, since the circuit is designed to work with genetic initial states, a different set of computational basis

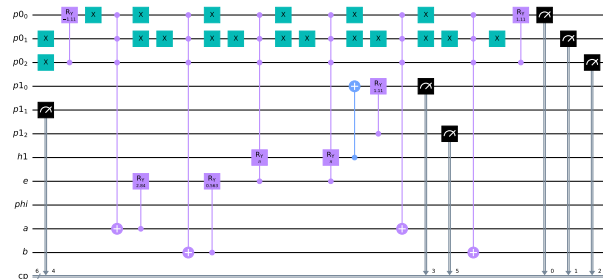


Figure 2: Quantum circuit for the $N_{\text{evol}} = 1$ step parton shower algorithm when the initial state is $|f_1\rangle$. Coupling constants are $g_1 = 2$ and $g_2 = g_{12} = 1$.

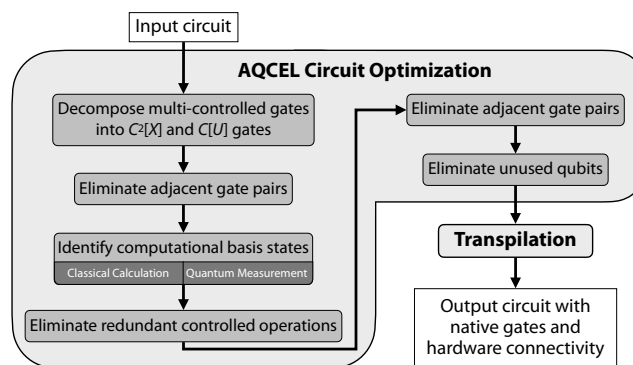


Figure 3: Flowchart of the proposed optimization protocol. We eliminate unnecessary gates, qubit controls and unused qubits. Finally, the resulting circuit can be encoded into particular gates for specific hardware.

states is occupied for each initial state, resulting in redundant controlled operations in the circuit.

The quantum circuit for $N_{\text{evol}} = 1$ step and the initial state $|f_1\rangle$ provide a good benchmark for the AQCEL protocol. Coupling constants are set to $g_1 = 2$ and $g_2 = g_{12} = 1$. Figure 2 shows the benchmark quantum circuit for the QPS algorithm with $N_{\text{evol}} = 1$.

3.2 Experimental setup

The AQCEL protocol focuses on circuit optimization at the algorithmic level, instead of at the level of a specific implementation using native gates for a particular quantum device. In addition to the initial-state dependent reduction of unnecessary controlled operations (Sec. 2.3 and 2.4), the AQCEL performs sequential decompositions of multi-controlled gates as well as the removal of adjacent gate pairs and unused qubits. A high-level flowchart of the AQCEL protocol is shown in Fig. 3.

The AQCEL is implemented using IBM Qiskit

version 0.32.1 [29] with Terra 0.18.3, Aer 0.9.1 and Ignis 0.6.0 APIs in Python 3.8.1 [68]. The codes and experimental results are available in GitHub [69]. We attempt to optimize circuits running on a classical computer with a single 2.4 GHz Intel core i5 processor. The AQCEL optimization performance is evaluated using the 27-qubit IBM device called *ibm_kawasaki* equipped with the IBM Quantum Falcon Processor and the statevector simulator in Qiskit Aer. When executing a circuit on the *ibm_kawasaki*, the gates in the circuit are transformed into machine-native single- (X, S_x, R_z) and two-qubit (CNOT) gates, and the qubits are mapped to the hardware, accounting for the actual qubit connectivity of the *ibm_kawasaki*. 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) and the simulator does not consider any quantum noise.

In addition to AQCEL, two circuit optimization tools are used: `t|ket)` in `pytket 0.17.0` and `pytket-qiskit 0.20.0`, and IBM Qiskit transpiler. They are used as references for the comparison of the optimization performance as well as in the combination with AQCEL to further reduce the gate counts. For `t|ket)`, the `get_compiled_circuit` routine with optimization level 2 is used ⁷. For Qiskit, the transpilation with optimization level 3 pass manager is applied ⁸. The decomposition of multi-controlled gates using relative phase Toffoli gates is applied first to all the cases for comparison on an equal footing.

3.3 Results

Here we discuss results for the AQCEL optimization to the $N_{\text{evol}} = 1$ QPS circuit. The figure of merit is the numbers of single-qubit gates (S_X, X) ⁹ and CNOT gates obtained by decomposing quantum gates in the circuit and the calculation time of the circuit. The calculation time is defined as the duration of the pulse schedule of the

⁷This routine is documented in the `pytket` manual at <https://cqcl.github.io/tket/pytket/api/backends.html>. The routine for decomposition all gates into basis gates is always applied in front of `t|ket)`.

⁸The value of `seed_transpiler` is fixed to 1 in order to suppress the randomness of Qiskit transpiler.

⁹ R_z gates are not included because they can be implemented with no cost as virtual Z gates [70].

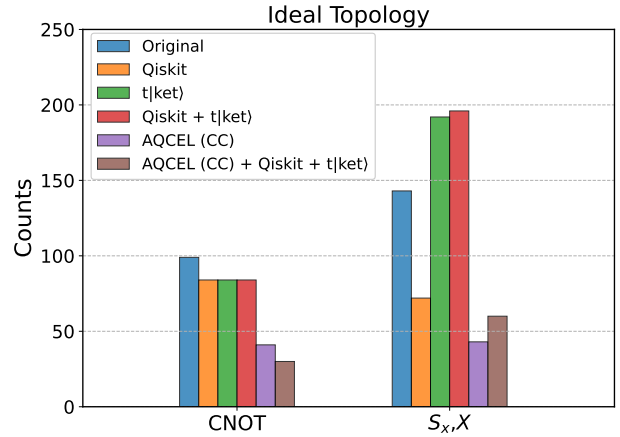


Figure 4: Numbers of CNOT gates and single-qubit gates of the QPS circuit transpiled considering the ideal topology. For the AQCEL optimization, the bitstring determination at controlled gates is performed using classical calculation.

transpiled circuit from the input to the measurements, as implemented in Qiskit.

First, we examine the numbers of single-qubit gates and CNOT gates assuming an ideal topology before and after the AQCEL optimization alone. The determination of bitstrings at controlled gates is performed using classical calculation. Figure 4 shows the gate counts from the original circuit and the circuits optimized using either AQCEL, `t|ket)` or Qiskit and their different combinations. It is seen that the AQCEL alone reduces gate counts drastically, and even more for CNOT gate when the AQCEL is combined with the `t|ket)` and Qiskit.

Starting with 155 CNOT and 431 single-qubit gates, the AQCEL alone removes 114 CNOT and 388 single-qubit gates, in which 58 CNOT and 88 single-qubit gates are accounted for by the reduction of redundant qubit controls ¹⁰, and the rest by the removal of adjacent gate pairs. The number of qubits is reduced from 14 to 13. The register n_ϕ , composed of only one qubit, is removed because it is used only for the case where the initial state is $|\phi\rangle$. The entire AQCEL optimization takes about 0.58 seconds. The wall time is by far dominated by the elimination of redundant controlled operations (that accounts for 65% of the total time), followed by a sub-dominant contribu-

¹⁰Removing redundant qubit controls reduces the cost of $C^2[X]$ and $C[U]$ decomposition into basis gates, resulting in the overall reduction of CNOT and single-qubit gate counts.

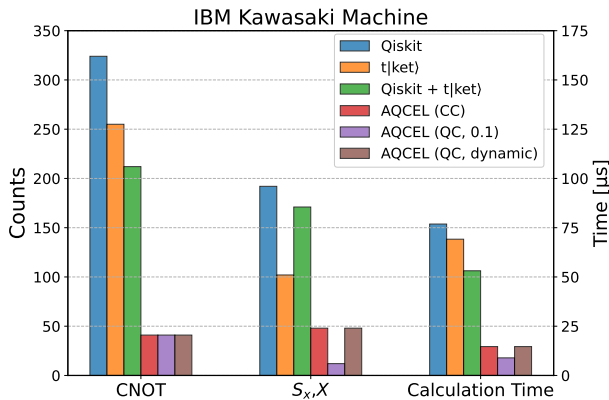


Figure 5: Numbers of gates and the calculation times of the QPS circuit transpiled considering the topology of *ibm_kawasaki*. For the AQCEL optimization, the bitstring determination at controlled gates is performed using either classical calculation (denoted by CC) or measuring the computational basis states on quantum hardware (denoted by QC) with the static thresholds of 0.1 or the dynamic threshold of s_ϵ^{dyn} .

tion of 35% from the adjacent gate-pair elimination.

Now we evaluate the performance of the optimizers with the transpilation considering the hardware topology of the *ibm_kawasaki*. In AQCEL, the bitstring determination is performed using classical calculation (denoted by CC) or quantum hardware (denoted by QC) with several thresholds. Figure 5 shows the results for the $N_{\text{evol}} = 1$ QPS circuit. The AQCEL reduces the CNOT gate counts more significantly than the case for ideal topology (Fig. 4). This reduction comes from less SWAP gates (each of which requires 3 CNOTs) for the AQCEL circuit because, once redundant qubit controls are removed, the amount of SWAP operations between multiple qubits is also suppressed. This results in much shorter calculation time for the AQCEL circuit. For the rest in the paper, the most efficient transpilation that combines the AQCEL, t|ket) and Qiskit is used for identification of the bitstrings at controlled gates and the fidelity measurement. The qubit counts are reduced from 14 to 13 under the dynamic threshold of s_ϵ^{dyn} . Under the static thresholds, the qubit counts reduce to 13 for $0.005 \leq s_\epsilon^f \leq 0.25$, but more significantly to 8 when $s_\epsilon^f = 0.3$.

In the initial-state dependent circuit optimization, what is preserved is the equivalence of the final state, not the circuit itself. To evaluate the accuracy of the AQCEL optimization, we con-

sider a classical fidelity between final states before and after the optimization, defined in terms of the probability distributions of the bitstrings observed in the measurement at the end of the circuits. This quantity, denoted as F and referred to as just “fidelity” hereafter, is given by

$$F = \sum_k \sqrt{p_k^{\text{orig}} p_k^{\text{opt}}}, \quad (18)$$

where the index k runs over the bitstrings. 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. For the ideal final state where any quantum error is not considered, the first fidelity, denoted F_{sim} , aims to quantify the amount of modifications to the final state introduced by the optimization procedure at the algorithmic level. To calculate the F_{sim} , both p^{orig} and p^{opt} are computed using the statevector simulation. The value of $F_{\text{sim}} = 1$ indicates that the final states are identical before and after the optimization (up to a possible phase difference on each of the qubits), while a deviation from unity gives a measure of how much the optimization has modified the final state from the ideal one.

The second fidelity value, F_{meas} , is computed using measurements with an actual quantum computer for p^{opt} . The p^{opt} is estimated from the rate at which a bitstring occurs in a large number of repeated measurements. The p^{orig} is computed using simulation, as for the F_{sim} . Even if $F_{\text{sim}} = 1$, the presence of noise will make $F_{\text{meas}} < 1$, with the difference from unity getting larger when more gates (particularly CNOT gates) are present in the circuit. Removing CNOT gates to optimize the circuit will lower the overall effect of noise and raise the F_{meas} value. However, when low-amplitude computational basis states are rejected by the threshold, more qubit controls are removed, which makes the final state different from the ideal one and decreases the F_{meas} value. Thus, the F_{meas} is a measure that reflects the trade-off of making the circuit shorter and changing the final state through the optimization. The measurements are performed 10000 times for each optimized circuit to obtain the F_{meas} value, and the experiment is repeated 30 times with the same optimized circuit to finally obtain the average and the standard deviation of the F_{meas}

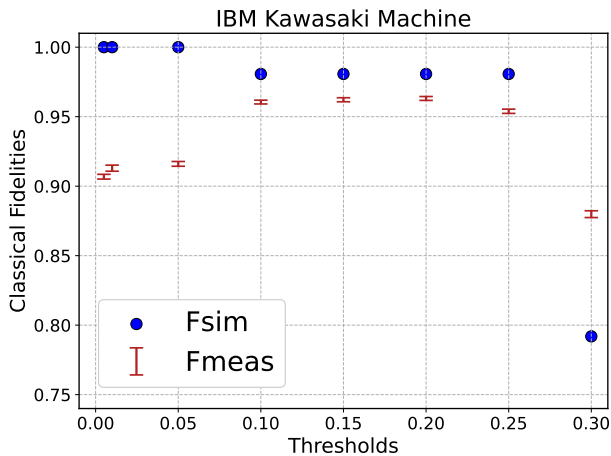


Figure 6: F_{sim} and F_{meas} values as a function of the static threshold s_{ϵ}^f used to optimize the QPS circuits by AQCEL. The bitstring measurement on controlled qubits is performed on the *ibm_kawasaki*.

values. Any error mitigation is not used in the measurements.

When the elimination of redundant qubit controls is performed based on measurements using a quantum computer with the static thresholds s_{ϵ}^f , the threshold dependence of F_{sim} and F_{meas} values is shown in Fig. 6. With increasing s_{ϵ}^f value the F_{meas} first increases, indicating the suppression of noise effects due to CNOT gate removal, but then worsens significantly at $s_{\epsilon}^f = 0.30$. This is understood from the behavior of the F_{sim} value: the F_{sim} stays close to unity up to $s_{\epsilon}^f = 0.25$ then decreases significantly, signaling that the optimization is too aggressive to maintain the final state from the original one if $s_{\epsilon}^f = 0.30$ is used. For the circuit considered here, the performance of the optimization appears to be best with $0.10 \leq s_{\epsilon}^f \leq 0.25$.

Shown in Fig. 7 is the F_{meas} versus CNOT gate counts with the determination of bitstrings using classical calculation and hardware measurement. For the $s_{\epsilon}^{\text{dyn}}$, the optimized circuit turns out to be exactly same as the one obtained using classical calculation¹¹. For the AQCEL optimized circuits, the F_{meas} values are 0.961 ± 0.001 for $s_{\epsilon}^f = 0.1$ and 0.911 ± 0.002 for the optimization using classical calculation. This demonstrates a clear improvement from the hardware measurement by removing qubit controls that trigger on

¹¹Although these circuits are identical, the F_{meas} values are slightly different due to statistical uncertainty and the quantum noise of the actual device.

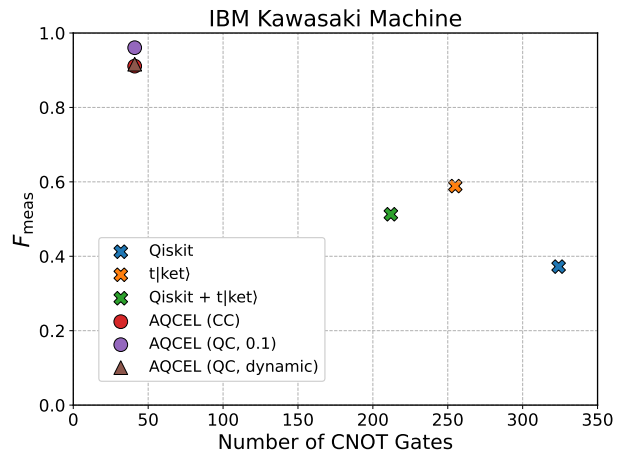


Figure 7: Fidelity F_{meas} versus the number of CNOT gates for the QPS circuit transpiled considering the topology of *ibm_kawasaki*. For the AQCEL optimization, the bitstring determination at controlled gates is performed using classical calculation (denoted by CC) or measuring the computational basis states on quantum hardware (denoted by QC) with the static thresholds of 0.1 or the dynamic threshold of $s_{\epsilon}^{\text{dyn}}$.

low-amplitude computational basis states. The F_{sim} value is 0.981 for $s_{\epsilon}^f = 0.1$, meaning that the final state is modified slightly from the ideal one.

To further evaluate the accuracy of the AQCEL optimized circuit, a quantum state tomography (QST) is performed over the particle registers of six qubits in the $N_{\text{evol}} = 1$ QPS circuit (see Fig. 2). From the measurements of 3^6 circuits, each performed 4000 times, with Pauli $\{X, Y, Z\}$ observables using the quantum hardware, the density matrix is reconstructed for the optimized circuit and compared with one obtained from the original circuit to compute a fidelity (denoted by F_{QST}). In a statevector simulation, the F_{QST} value is unity for both Qiskit+t|ket) and AQCEL circuit with the $s_{\epsilon}^{\text{dyn}}$, meaning that the optimization does not modify the final state including relative phases. The F_{QST} values are measured to be 0.13 and 0.41 for Qiskit+t|ket) and AQCEL circuits, respectively, when performing QST on the quantum hardware. The F_{QST} values are much smaller than unity due to hardware noise, but the AQCEL shows a clear improvement over the Qiskit+t|ket).

4 Discussion

4.1 Applicability of proposed optimization

The core component of the proposed circuit optimization is the identification of computational basis states with zero- or low-amplitudes using a quantum computer and the subsequent elimination of redundant controlled operations. Therefore, the 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, the AQCEL would not be effective if all the computational basis states have non-negligible amplitudes, especially when they are small in size because of thresholds applied on quantum hardware. A typical example is Quantum Phase Estimation [71] or Grover's Algorithm, where an equal superposition state is created first by applying $H^{\otimes n}$ gates to the initial state $|0\rangle^{\otimes n}$ of the n -qubit system. However, even in this case, the AQCEL can efficiently produce a quantum circuit that approximates the original final state by ignoring low-amplitude computational basis states.

Another important aspect for the AQCEL optimization protocol is the resource needed to use a quantum computer for the optimization. In the NISQ era, it is worth spending the quantum computer resource in the optimization if the resulting circuit can produce higher-fidelity results than the original circuit does. In the fault-tolerant quantum computing era, using a quantum computer for the optimization may not be crucial due to its capability of correcting quantum errors during the operation. However, the initial-state dependent AQCEL optimization will be still useful for simplifying the quantum circuit, even in the fault-tolerant regime.

4.2 Further simplifications with initial-state dependent optimization

The new method proposed here for obtaining bit-strings using a quantum computer will open new possibilities in initial-state dependent quantum circuit optimization within polynomial computational resources. In this paper, we discuss the simplest example of this type of optimizations, that is, just use the information of control qubits and optimize multi-controlled qubit gates individually. There are several possibilities to extend

the idea for further simplifications, e.g., using ancilla qubits, quantum gates with qutrit (3-level) states, or adding the information of target qubit in controlled gate like RPO [49]. One can optimize not only individual gates but also multiple gates as a gate set in future.

As another interesting possibility, if a circuit turns out to contain only a small number of basis states, one could represent the circuit state using fewer qubits than the original ones. Given that this approach might require a completely new computational basis, this is left for future work.

4.3 Mitigations of quantum errors

The threshold choice in AQCEL has significant impacts on F_{sim} and F_{meas} , as seen in Figs. 6 and 7. The measurement error can be improved by adapting the unfolding technique developed in Ref. [17] and related approaches that use fewer resources [20, 72, 73, 74, 75] or further mitigate the errors [22]. 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 with identity insertions, first proposed in Ref. [23] and generalized in Ref. [27]. The method based on the CNOT error mitigation may improve the accuracy of optimizations and the fidelity of our approach.

5 Conclusion and outlook

We have proposed a new optimization protocol, called AQCEL, for analyzing quantum circuits to remove redundant controlled operations. The AQCEL can remove unnecessary qubit controls from multi-controlled gates even when all the relevant qubits are entangled. 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 that trigger on low-amplitude basis states will result in a circuit that produces the final state distinct from the original one, this may be a desirable feature un-

der the existence of hardware noise.

We have adopted the AQCEL optimization scheme to the quantum parton shower simulation using the *ibm_kawasaki*. In the experiment, the proposed scheme has shown a significant reduction of gate counts and improved the F_{meas} value while retaining the accuracy of the probability distributions of the final state. For the AQCEL optimized circuits, the F_{meas} values are 0.961 ± 0.001 for the static threshold $s_\epsilon^f = 0.1$ and 0.911 ± 0.002 for the optimization using classical calculation. For the dynamic threshold s_ϵ^{dyn} , the optimized circuit is exactly same as the one obtained using classical calculation. The initial-state dependent optimization discussed here opens new possibilities to extend quantum circuit optimization further in future.

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.

This study is partly carried out under the project ‘‘Optimization of HEP Quantum Algorithms’’ supported by the U.S.-Japan Science and Technology Cooperation Program in High Energy Physics.

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 Decomposition of multi-controlled gates

When the same qubit controls are removed from $C^m[U]$ and $C^s[U]$ ($m > s$), Eq. (11) indicates that the condition for the removal of qubit controls is stricter for the latter because the number of k is larger. This means that if we decompose a $C^m[U]$ into controlled gates with smaller number of control qubits, e.g., Toffoli and two-qubit gates, the opportunity to remove redundant qubit controls

is partly lost. This suggests that removing redundant qubit controls should be applied before decomposing multi-controlled gates into basis gates like CNOT. As mentioned in Sec. 2.2 and Appendix B, the decomposition of multi-controlled gates is essential in terms of computational resources. This means that there is a trade-off between the opportunities of removing qubit controls and computational complexity.

B Computational resources for the proposed optimization scheme

The computational costs to perform the proposed optimization scheme are 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 first step in the optimization scheme is the identification of computational basis states. If we use the classical calculation to simply track all the computational basis states whose amplitudes may be nonzero at each point of the circuit, it requires the computation of $\mathcal{O}(N2^n)$ states which grows exponentially with n . This method requires less computational resource than a statevector simulation, but it neglects certain rare cases where exact combinations of amplitudes lead to the elimination of redundant controlled operations. This is because a statevector simulation is applied as a classical calculation for the identification of bitstrings in AQCEL. 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

$$\begin{aligned} M\{m + (1 + m) + (2 + m) + \cdots + (N - 1 + m)\} \\ = \frac{1}{2}MN(N - 1) + mMN \quad (19) \end{aligned}$$

Therefore, the computational cost grows polynomially with n in $\mathcal{O}(MN^2 + mMN)$.

We next consider the decision of all 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 bitstrings is 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. We need to check

if all specified bitstrings in Eq. (11) for the chosen combination are not measured. The cost is $\mathcal{O}(Mm2^m)$ for one chosen combination because the size of the bitstring is m and the numbers of measured and specified bitstrings are M and 2^m , respectively. Therefore, the overall computational cost for the determination of redundant qubit controls is $\mathcal{O}(Mm4^mN)$ for N multi-qubit controlled gates. The classical calculation requires $\mathcal{O}(m8^mN)$ as well.

In the AQCEL protocol, all controlled gates in the circuit are decomposed into Toffoli gates and two-qubit controlled- U gates. 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 redundant qubit control identification becomes polynomial in $\mathcal{O}(MmN)$ because all controlled gates have the constant number of control qubits ($m = 1, 2$). The computational cost for the identification of computational basis states still behaves polynomially in $\mathcal{O}(m^2MN^2)$ when a quantum computer is used. Given that a controlled gate has at most $n - 1$ control qubits, the total computational cost for the entire optimization sequence is $\mathcal{O}(n^2MN^2)$, when the computational basis state measurement is performed using a quantum computer.

References

- [1] John Preskill. “Quantum Computing in the NISQ era and beyond”. *Quantum* **2**, 79 (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**, 375–379 (2017).
- [3] Alexander Zlokapa, Alex Mott, Joshua Job, Jean-Roch Vlimant, Daniel Lidar, and Maria Spiropulu. “Quantum adiabatic machine learning by zooming into a region of the energy surface”. *Phys. Rev. A* **102**, 062405 (2020).
- [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 Lepton-Photon2019*, 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**, 2 (2021).
- [6] Wen Guan, Gabriel Perdue, Arthur Pesah, Maria Schuld, Koji Terashi, Sofia Vallecorsa, and Jean-Roch Vlimant. “Quantum machine learning in high energy physics”. *Mach. Learn.: Sci. Technol.* **2**, 011003 (2021).
- [7] Vasilis Belis, Samuel González-Castillo, Christina Reissel, Sofia Vallecorsa, Elías F. Combarro, Günther Dissertori, and Florentin Reiter. “Higgs analysis with quantum classifiers”. *EPJ Web Conf.* **251**, 03070 (2021).
- [8] 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”. *Quantum Mach. Intell.* **3**, 27 (2021).
- [9] 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”. *EPJ Web Conf.* **245**, 09013 (2020).
- [10] Ilyya Shapoval and Paolo Calafiura. “Quantum Associative Memory in HEP Track Pattern Recognition”. *EPJ Web Conf.* **214**, 01012 (2019).
- [11] 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 (2020).
- [12] Annie Y. Wei, Preksha Naik, Aram W. Harrow, and Jesse Thaler. “Quantum algorithms for jet clustering”. *Phys. Rev. D* **101**, 094015 (2020).
- [13] Souvik Das, Andrew J. Wildridge, Sachin B. Vaidya, and Andreas Jung. “Track clustering with a quantum annealer for primary vertex reconstruction at hadron col-

- leaders”. [arXiv:1903.08879 \[hep-ex\]](#) (2019) [arXiv:1903.08879](#).
- [14] Kyle Cormier, Riccardo Di Sipio, and Peter Wittek. “Unfolding measurement distributions via quantum annealing”. *JHEP* **11**, 128 (2019).
- [15] Davide Provasoli, Benjamin Nachman, Christian Bauer, and Wibe A de Jong. “A quantum algorithm to efficiently sample from interfering binary trees”. *Quantum Sci. Technol.* **5**, 035004 (2020).
- [16] Benjamin Nachman, Davide Provasoli, Wibe A. de Jong, and Christian W. Bauer. “Quantum algorithm for high energy physics simulations”. *Phys. Rev. Lett.* **126** (2021).
- [17] Christian W. Bauer, Wibe A. De Jong, Benjamin Nachman, and Miroslav Urbanek. “Unfolding quantum computer readout noise”. *npj Quantum Inf.* **6**, 84 (2020).
- [18] Yanzhu Chen, Maziar Farahzad, Shinjae Yoo, and Tzu-Chieh Wei. “Detector tomography on IBM 5-qubit quantum computers and mitigation of imperfect measurement”. *Phys. Rev. A* **100**, 052315 (2019).
- [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”. *Phys. Rev. Lett.* **108**, 057002 (2012).
- [20] Michael R Geller and Mingyu Sun. “Toward efficient correction of multiqubit measurement errors: pair correlation method”. *Quantum Sci. Technol.* **6**, 025009 (2021).
- [21] Michael R Geller. “Rigorous measurement error correction”. *Quantum Sci. Technol.* **5**, 03LT01 (2020).
- [22] Rebecca Hicks, Christian W. Bauer, and Benjamin Nachman. “Readout rebalancing for near-term quantum computers”. *Phys. Rev. A* **103**, 022407 (2021).
- [23] 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 (2018).
- [24] Suguru Endo, Simon C. Benjamin, and Ying Li. “Practical Quantum Error Mitigation for Near-Future Applications”. *Phys. Rev. X* **8**, 031027 (2018).
- [25] Kristan Temme, Sergey Bravyi, and Jay M. Gambetta. “Error Mitigation for Short-Depth Quantum Circuits”. *Phys. Rev. Lett.* **119**, 180509 (2017).
- [26] 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**, 491–495 (2019).
- [27] 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 (2020).
- [28] Matthew Otten and Stephen K. Gray. “Recovering noise-free quantum observables”. *Phys. Rev. A* **99**, 012338 (2019).
- [29] Gadi Aleksandrowicz, *et al.* “Qiskit: An Open-source Framework for Quantum Computing”. *Zenodo*. (2019).
- [30] Seyon Sivarajah, Silas Dilkes, Alexander Cowtan, Will Simmons, Alec Edgington, and Ross Duncan. “t|ket): a retargetable compiler for NISQ devices”. *Quantum Sci. Technol.* **6**, 014003 (2020).
- [31] Thomas Häner, Damian S Steiger, Krysta Svore, and Matthias Troyer. “A software methodology for compiling quantum programs”. *Quantum Sci. Technol.* **3**, 020501 (2018).
- [32] Alexander S. Green, Peter LeFanu Lumsdaine, Neil J. Ross, Peter Selinger, and Benoît Valiron. “Quipper: a scalable quantum programming language”. *SIGPLAN Not.* **48**, 333–342 (2013).
- [33] 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 Comput.* **45**, 2–17 (2015).
- [34] Krysta Svore, Martin Roetteler, Alan Geller, Matthias Troyer, John Azariah, Christopher Granade, Bettina Heim, Vadym Kliuchnikov, Mariia Mykhailova, and Andres Paz. “Q#: Enabling Scalable Quantum Computing and Development with a High-level DSL”. In *Proceedings of the Real World Domain Specific Languages Workshop 2018*.

- Pages 1–10. Association for Computing Machinery (2018).
- [35] 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 (2019).
- [36] Robert S. Smith, Michael J. Curtis, and William J. Zeng. “A Practical Quantum Instruction Set Architecture”. arXiv:1608.03355 [quant-ph] (2016) arXiv:1608.03355.
- [37] Damian S. Steiger, Thomas Häner, and Matthias Troyer. “ProjectQ: an open source software framework for quantum computing”. *Quantum* **2**, 49 (2018).
- [38] Cirq Developers. “Cirq”. *Zenodo*. (2021).
- [39] Alexander J. McCaskey, Eugene F. Dumitrescu, Dmitry Liakh, Mengsu Chen, Wuchun Feng, and Travis S. Humble. “A language and hardware independent approach to quantum–classical computing”. *SoftwareX* **7**, 245–254 (2018).
- [40] 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”. In Proceedings of the 46th International Symposium on Computer Architecture. Pages 527–540. (2019).
- [41] Robert S Smith, Eric C Peterson, Erik J Davis, and Mark G Skilbeck. “quilc: An Optimizing Quil Compiler”. *Zenodo*. (2020).
- [42] Yunseong Nam, Neil J. Ross, Yuan Su, Andrew M. Childs, and Dmitri Maslov. “Automated optimization of large quantum circuits with continuous parameters”. *npj Quantum Inf.* **4**, 23 (2018).
- [43] 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). url: api.semanticscholar.org/CorpusID:115143379.
- [44] Prakash Murali, Jonathan M. Baker, Ali Javadi Abhari, Frederic T. Chong, and Margaret Martonosi. “Noise-Adaptive Compiler Mappings for Noisy Intermediate-Scale Quantum Computers”. arXiv:1901.11054 [quant-ph] (2019) arXiv:1901.11054.
- [45] Prakash Murali, David C. McKay, Margaret Martonosi, and Ali Javadi-Abhari. “Software Mitigation of Crosstalk on Noisy Intermediate-Scale Quantum Computers”. In Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems. Pages 1001–1016. ASPLOS ’20. Association for Computing Machinery (2020).
- [46] Eric C. Peterson, Gavin E. Crooks, and Robert S. Smith. “Fixed-Depth Two-Qubit Circuits and the Monodromy Polytope”. *Quantum* **4**, 247 (2020).
- [47] Nelson Leung, Mohamed Abdelhafez, Jens Koch, and David Schuster. “Speedup for quantum optimal control from automatic differentiation based on graphics processing units”. *Phys. Rev. A* **95**, 042318 (2017).
- [48] 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”. In Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture. Page 266–278. Association for Computing Machinery (2019).
- [49] Ji Liu, Luciano Bello, and Huiyang Zhou. “Relaxed Peephole Optimization: A Novel Compiler Optimization for Quantum Circuits”. arXiv:2012.07711 [quant-ph] (2020) arXiv:2012.07711.
- [50] 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 (1995).
- [51] Dmitri Maslov. “Advantages of using relative-phase Toffoli gates with an application to multiple control Toffoli optimization”. *Phys. Rev. A* **93**, 022311 (2016).
- [52] D. Michael Miller, Robert Wille, and Rolf Drechsler. “Reducing reversible circuit cost by adding lines”. In 2010 40th IEEE Interna-

- tional Symposium on Multiple-Valued Logic. Pages 217–222. (2010).
- [53] Pranav Gokhale, Jonathan M. Baker, Casey Duckering, Natalie C. Brown, Kenneth R. Brown, and Frederic T. Chong. “Asymptotic improvements to quantum circuits via qutrits”. In Proceedings of the 46th International Symposium on Computer Architecture. Page 554–566. Association for Computing Machinery (2019).
- [54] Yushi Wang and Marek Perkowski. “Improved complexity of quantum oracles for ternary grover algorithm for graph coloring”. In 2011 41st IEEE International Symposium on Multiple-Valued Logic. Pages 294–301. (2011).
- [55] Alexey Galda, Michael Cubeddu, Naoki Kanazawa, Prineha Narang, and Nathan Earnest-Noble. “Implementing a Ternary Decomposition of the Toffoli Gate on Fixed-Frequency Transmon Qutrits”. arXiv:2109.00558 [quant-ph] (2021) arXiv:2109.00558.
- [56] Toshiaki Inada, Wonho Jang, Yutaro Iiyama, Koji Terashi, Ryu Sawada, Junichi Tanaka, and Shoji Asai. “Measurement-Free Ultrafast Quantum Error Correction by Using Multi-Controlled Gates in Higher-Dimensional State Space”. arXiv:2109.00086 [quant-ph] (2021) arXiv:2109.00086.
- [57] Yuchen Wang, Zixuan Hu, Barry C. Sanders, and Sabre Kais. “Qudits and High-Dimensional Quantum Computing”. *Front. Phys.* **8**, 479 (2020).
- [58] T. C. Ralph, K. J. Resch, and A. Gilchrist. “Efficient Toffoli gates using qudits”. *Phys. Rev. A* **75**, 022313 (2007).
- [59] E. O. Kiktenko, A. S. Nikolaeva, Peng Xu, G. V. Shlyapnikov, and A. K. Fedorov. “Scalable quantum computing with qudits on a graph”. *Phys. Rev. A* **101**, 022304 (2020).
- [60] Jing Zhong and Jon C. Muzio. “Using Crosspoint Faults in Simplifying Toffoli Networks”. In 2006 IEEE North-East Workshop on Circuits and Systems. Pages 129–132. (2006).
- [61] Ketan N. Patel, Igor L. Markov, and John P. Hayes. “Optimal synthesis of linear reversible circuits”. *Quantum Inf. Comput.* **8**, 282–294 (2008).
- [62] Matthew Amy, Parsiad Azimzadeh, and Michele Mosca. “On the controlled-NOT complexity of controlled-NOT–phase circuits”. *Quantum Sci. Technol.* **4**, 015002 (2018).
- [63] Sumeet Khatri, Ryan LaRose, Alexander Poremba, Lukasz Cincio, Andrew T. Sornborger, and Patrick J. Coles. “Quantum-assisted quantum compiling”. *Quantum* **3**, 140 (2019).
- [64] Tyson Jones and Simon C. Benjamin. “Robust quantum compilation and circuit optimisation via energy minimisation”. *Quantum* **6**, 628 (2022).
- [65] Bob Coecke and Ross Duncan. “Interacting quantum observables: categorical algebra and diagrammatics”. *New J. Phys.* **13**, 043016 (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 (2020).
- [67] Miriam Backens. “The ZX-calculus is complete for stabilizer quantum mechanics”. *New J. Phys.* **16**, 093021 (2014).
- [68] Guido Van Rossum and Fred L. Drake. “Python 3 Reference Manual”. CreateSpace. Scotts Valley, CA (2009). url: dl.acm.org/doi/book/10.5555/1593511.
- [69] UTokyo-ICEPP. “AQCEL”. GitHub. (2022). url: github.com/UTokyo-ICEPP/aqcel.
- [70] David C. McKay, Christopher J. Wood, Sarah Sheldon, Jerry M. Chow, and Jay M. Gambetta. “Efficient Z gates for quantum computing”. *Phys. Rev. A* **96**, 022330 (2017).
- [71] Michael A. Nielsen and Isaac L. Chuang. “Quantum Computation and Quantum Information”. Cambridge University Press. (2000).
- [72] 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”. *Phys. Rev. Lett.* **119**, 180511 (2017).
- [73] 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”. *Phys. Rev. Lett.* **122**, 110501 (2019).

- [74] 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 Greenberger-Horne-Zeilinger states via multiple quantum coherences”. *Phys. Rev. A* **101**, 032343 (2020).
- [75] Kathleen E. Hamilton, Tyler Kharazi, Titus Morris, Alexander J. McCaskey, Ryan S. Bennink, and Raphael C. Pooser. “Scalable quantum processor noise characterization”. arXiv:2006.01805 [quant-ph] (2020) arXiv:2006.01805.