

# Clifford recompilation for faster classical simulation of quantum circuits

Hamam Qassim<sup>1,3</sup>, Joel J. Wallman<sup>2,3</sup>, and Joseph Emerson<sup>1,2,3</sup>

<sup>1</sup>Department of Physics and Astronomy, University of Waterloo, Waterloo, Canada

<sup>2</sup>Department of Applied Mathematics, University of Waterloo, Waterloo, Canada

<sup>3</sup>Institute for Quantum Computing, Waterloo, Canada

February 8, 2019

Simulating quantum circuits classically is an important area of research in quantum information, with applications in computational complexity and validation of quantum devices. One of the state-of-the-art simulators, that of Bravyi et al, utilizes a randomized sparsification technique to approximate the output state of a quantum circuit by a stabilizer sum with a reduced number of terms. In this paper, we describe an improved Monte Carlo algorithm for performing randomized sparsification. This algorithm reduces the runtime of computing the approximate state by the factor  $\ell/m$ , where  $\ell$  and  $m$  are respectively the total and non-Clifford gate counts. The main technique is a circuit recompilation routine based on manipulating exponentiated Pauli operators. The recompilation routine also facilitates numerical search for Clifford decompositions of products of gates, which can further reduce the runtime in certain cases. It may additionally lead to a framework for optimizing circuit implementations over a gate-set, reducing the overhead for state-injection in fault-tolerant implementations. We provide a concise exposition of randomized sparsification, and describe how to use it to estimate circuit amplitudes in a way which can be generalized to a broader class of gates and states. This latter method can be used to obtain additive error estimates of circuit probabilities with a faster runtime than the full techniques of Bravyi et al. Such estimates are useful for validating near-term quantum devices provided that the target probability is not exponentially small.

## 1 Introduction

Digital simulation of quantum phenomena is a central problem in computational physics, with deep theoretical and practical significance. On the theory side, research into this problem has produced surprising connections between different branches of mathematics and physics, while providing major insight into the inner-workings of quantum theory. On the practical side, there is high demand for faster and more accurate simulations in fields such as quantum computing, materials science, and quantum chemistry. In these disciplines, it is desirable to have a toolkit of simulation techniques which can handle the wide variety of quantum systems encountered.

Quantum computers are widely believed to be hard to simulate classically. Nevertheless, this hardness is only asymptotic, and it is always possible in principle to simulate larger and larger quantum computers by

Hamam Qassim: [hqassim@uwaterloo.ca](mailto:hqassim@uwaterloo.ca)

optimizing the usage of classical resources. Efforts to demonstrate such optimizations have surged recently [2, 6–8, 15]. These efforts are partially motivated by upcoming quantum computers containing on the order of 100 qubits [13, 14]. Since these prototypes are likely to be considerably noisy, the ability to simulate their ideal counterparts is essential for validating their performance and to aid in their design and debugging. Furthermore, simulation of quantum circuits is directly tied to establishing when quantum computers are useful, as demonstrating quantum supremacy requires sampling from distributions which are provably hard to (approximately) sample from using classical means [10].

Near-Clifford circuits constitute an important class of quantum circuits. Fault-tolerant implementations of quantum computing in the near-term are likely to contain a small number of non-Clifford gates. This is the case since the overhead associated with implementing non-Clifford gates fault-tolerantly is huge, for example accounting for nearly 90% of the number of physical

qubits in surface code implementations of Shor’s algorithm [9]. A host of classical simulation schemes have recently been proposed for simulating near-Clifford circuits, with runtimes exhibiting mild exponential dependence on  $n$  [1, 3–5, 12]. State-of-the-art techniques rely on decomposing the  $n$ -qubit output state of a near-Clifford circuit  $U$  as a superposition of stabilizer states [3, 5],

$$|\psi\rangle \equiv U|0^n\rangle = \sum_{i=1}^N a_i |v_i\rangle, \quad (1)$$

with  $N \ll 2^n$ . These techniques allow for approximate sampling from the outcome distribution of the circuit, with the runtime for producing a single sample scaling like  $O(N\epsilon^{-2})$ , where  $\epsilon$  is the error in total variation distance [3, 5]. The minimum number of terms over all such decompositions of  $|\psi\rangle$  is called its *stabilizer rank*, and is a measure of the complexity of simulating  $|\psi\rangle$ . Finding stabilizer decompositions with a small number of terms is typically quite hard, even for simple product states [3, 5]. This motivates using *approximate* stabilizer decompositions. It was shown in [5] that an approximation to  $|\psi\rangle$  with a smaller number of stabilizer terms can be obtained by “randomly sparsifying” any decomposition of the type in eq. (1). This is done by generating  $k$  random stabilizer states, each chosen independently to be  $|v_i\rangle$  with probability  $|a_i|/\|a\|_1$ , then setting the approximate state  $|\tilde{\psi}\rangle$  to be their equal superposition (with a uniform scaling factor). If  $k$  is large enough (roughly  $\|a\|_1^2 \epsilon^{-2}$ ), then  $\|\psi - \tilde{\psi}\| \leq \epsilon$  with high probability. Typically  $k \ll N$ , so that using the approximate state  $|\tilde{\psi}\rangle$  for simulation purposes is preferable.

Randomized sparsification can be performed using a sum-over-Cliffords Monte Carlo method [5], which we briefly review before stating our contribution. Given a circuit  $U$ , each non-Clifford gate in  $U$  can be expanded as a linear combination of Clifford gates. The entire circuit can then be expressed as

$$U = \sum_i a_i C_i, \quad (2)$$

for a collection of (typically exponentially many) Clifford circuits  $C_i$ . The sum-over-Cliffords method uses this decomposition to compute the stabilizer states in eq. (1) while preserving the relative phases between them. For each term, a phase-sensitive Clifford simulator is used to convert the description of the circuit  $C_i$  to a description of the stabilizer state  $|v_i\rangle = C_i|0^n\rangle$  while keeping track of the global phase. Similarly, the stabilizer terms of the approximate state  $|\tilde{\psi}\rangle$  can be computed by applying circuit  $C_i$  with probability  $|a_i|/\|a\|_1$ . The Clifford simulator computes the so-called CH-form of  $C_i|0^n\rangle$  by performing an update for each gate in  $C_i$ .

This computation takes time  $O(\ell n^2)$ , where  $\ell$  is the total number of gates in the circuit. Thus to compute the CH-form for all stabilizer terms comprising the approximate state  $|\tilde{\psi}\rangle$ , the method in [5] runs in time  $O(k\ell n^2)$ .

In the current paper, we report an improved variant of randomized sparsification based on recompiling the circuit using exponentiated Pauli operators. This variant removes the dependence on the total number of gates in the circuit. More precisely, it reduces the time cost of computing the CH description of  $|v_i\rangle$  from  $O(\ell n^2)$  to  $O(mn^2)$ , where  $m$  is the non-Clifford gate count in  $U$ . This reduces the runtime of key subroutines in [5] by the factor  $\ell/m$ , which can be significant for circuits in which the number of Clifford gates is very large compared to non-Clifford gates. This is relevant, for example, in the case of the QAOA circuit simulated in [5], which contains hundreds of Clifford gates and only 64 T gates. The above savings are achieved at an additional memory cost of  $O(nm)$  bits, which is negligible in the regime of interest.

As a side benefit, the Clifford recompilation method makes it easier to numerically search for decompositions of the kind in eq. (1) with a reduced 1-norm. As the exponential scaling of the runtimes in [5] is determined by  $\|a\|_1^2$ , this allows for a significant speed-up in some cases. Furthermore, the total number of non-Clifford gates in the circuit can be reduced if adjacent gates in the recompiled circuit cancel out. This suggests a systematic way of optimizing the compilation step to reduce the number of non-Clifford gates in a given circuit, with the possibility of significant overhead savings in practical fault-tolerant implementations.

Another use of Clifford recompilation is in estimating circuit amplitudes. It is straight-forward to see that the random variable  $\langle x|\tilde{\psi}\rangle$  is an unbiased estimator of the amplitude  $\langle x|\psi\rangle$ , where  $x \in \mathbb{F}_2^n$  denotes a computational basis state. Clifford recompilation reduces the complexity of evaluating  $\langle x|\tilde{\psi}\rangle$  by the factor  $\ell/m$ , in comparison to [5]. Estimates of the amplitude obtained in this way are used in the Metropolis simulator in this latter work. They can also be used to estimate the probability  $|\langle x|\psi\rangle|^2$  up to additive error  $\epsilon$  with a runtime scaling like  $O(\|a\|_1^2 \epsilon^{-2})$ . Learning probabilities in this way can be generalized to a wider range of quantum circuits. It is also much less memory intensive, and is better suited for massive parallelization and GPU-computing. However, in general such polynomial precision estimates are only useful as long as the target probability is not exponentially small.

*Notation.* We write  $X, Y$ , and  $Z$  for the single qubit Pauli matrices, and typically use the letters  $P, Q$  and  $W$  to refer to Hermitian elements of the  $n$ -qubit Pauli group. The shorthand  $M_{b:a}$ , with  $b > a$ , denotes the product  $M_b M_{b-1} \dots M_a$ .

## 2 Improved randomized sparsification

**The CH-form of stabilizer states.** We begin by reviewing the CH-form of stabilizer states. A C-type Clifford is one which satisfies  $U_C|0^n\rangle = |0^n\rangle$ . An H-type Clifford consists of Hadamard gates on some subset of the qubits. Any stabilizer state of  $n$ -qubits can be written as

$$|\phi\rangle = \omega U_C U_H |s\rangle, \quad (3)$$

where  $U_C$  is a C-type Clifford,  $U_H$  is an H-type Clifford,  $s \in \mathbb{Z}_2^n$ , and  $\omega$  is a complex number.  $U_C$  can be described by a stabilizer tableau specifying its action on the Pauli operators  $X_i$  and  $Z_i$  for  $i = 1, \dots, n$ , and  $U_H$  can be described by an  $n$ -bit string specifying which qubits are acted on. As shown in [5], the tuple  $(\omega, U_C, U_H)$  describing the state  $|\phi\rangle$  can be updated in time  $O(n)$  under acting on  $|\phi\rangle$  with a gate from the set {CNOT, CPHASE, PHASE}, and in time  $O(n^2)$  under applying a Hadamard gate. For our circuit recompilation, we need the following fact concerning updates of the CH-form under Clifford gates of a particular exponential form.

**Proposition 1.** *The tuple  $(\omega, U_C, U_H)$  can be updated in time  $O(n^2)$  under applying a Clifford gate of the form  $\exp(i\theta P)$ , where  $P$  is a Pauli and  $\theta \in (\pi/4)\mathbb{Z}$ .*

*Proof.* Without loss of generality we assume that  $\theta \in \{\pi/2, \pi/4\}$ . We can compute a Pauli  $P'$  such that  $(U_C U_H)^\dagger P U_C U_H = P'$  in time  $O(n^2)$ , using the stabilizer tableau of  $U_C$  and the fact that  $HXH = Z$ . Given such  $P'$ , the state can be updated in time  $O(n^2)$ . This follows from the following observations. If  $\theta = \pi/2$ , then  $\exp(i\theta P)wU_C U_H |s\rangle = wU_C U_H (iP')|s\rangle = w'U_C U_H |s'\rangle$ , for some  $w'$  and  $s'$  which can be computed in time  $O(n)$ . If  $\theta = \pi/4$ , then  $\exp(i\theta P)wU_C U_H |s\rangle = wU_C U_H \exp(i\frac{\pi}{4}P')|s\rangle = w'U_C U_H (|s\rangle + i^\delta |s'\rangle)$ , for some  $\delta \in \mathbb{Z}_4$ ,  $w'$ , and  $s'$ , all of which can be computed in time  $O(n)$ . If  $s = s'$  then the update is trivial. If  $s \neq s'$ , then the update can be computed in time  $O(n^2)$  by combining Proposition 4 in [5] with the update rule for  $U_C$  under right-multiplication by gates from the set {CNOT, CPHASE, PHASE}, which is given in the same section in [5].  $\square$

**Clifford recompilation.** We now describe recompiling the circuit in order to reduce the complexity of computing the CH-forms of the output state. Consider a quantum circuit  $U$  consisting of  $\ell$  gates,  $m$  of which are non-Clifford gates. We can write  $U$  in the form

$$U = C_m U_m C_{m-1} U_{m-1} \dots C_1 U_1 C_0, \quad (4)$$

where  $C_i$  are Clifford circuits, and  $U_i$  are non-Clifford gates. We can then “pull” all non-Clifford gates to the end of the circuit

$$U = U'_m \dots U'_1 C_m C_{m-1} \dots C_1 C_0, \quad (5)$$

where  $U'_i = C_{m:i} U_i C_{m:i}^\dagger$ . We can absorb the action of the Clifford gates into the input state, so that the output state is given by

$$|\psi\rangle = U|0^n\rangle = U'|\phi\rangle, \quad U' \equiv U'_{m:1}, \quad (6)$$

where  $|\phi\rangle \equiv C_{m:0}|0^n\rangle$  is a stabilizer state. Note that the CH-form of  $|\phi\rangle$  can be computed in time  $O(cn^2)$ , where  $c \equiv \ell - m$  is the number of Clifford gates in  $U$ . Storing the CH-form of  $|\psi\rangle$  in memory requires  $O(n^2)$  bits.

Next we discuss how to compute the  $U'_i$ . Suppose we are given a Clifford decomposition of each non-Clifford gate in  $U$ ;

$$U_i = \sum_j b_{ij} V_{ij}, \quad (7)$$

where each  $V_{ij}$  is a product of a small number of elementary Clifford gates. We will use exponentiated Pauli matrices to compute and store  $U'_i$  efficiently. First note that each  $V_{ij}$  can be written as a product of Clifford gates of the form  $\exp(i\theta P)$  where  $P$  is a Pauli and  $\theta \in (\pi/4)\mathbb{Z}$ <sup>1</sup>. We can therefore write

$$U_i = \sum_j a_{ij} \prod_{k=1}^{O(1)} \exp(i\theta_{ijk} P_{ijk}), \quad (9)$$

where  $a_{ij}$  absorbs any phases from the elementary gates, and  $\{P_{ijk}\}$  is some collection of Pauli matrices. Then

$$\begin{aligned} U'_i &= \sum_j b_{ij} C_{m:i} \left[ \prod_k \exp(i\theta_{ijk} P_{ijk}) \right] C_{m:i}^\dagger \\ &= \sum_j b_{ij} \prod_k \exp(i\theta_{ijk} C_{m:i} P_{ijk} C_{m:i}^\dagger) \\ &= \sum_j b_{ij} \prod_k \exp(i\theta'_{ijk} P'_{ijk}) \\ &= \sum_j b_{ij} V'_{ij}. \end{aligned} \quad (10)$$

<sup>1</sup>For example, the Hadamard, PHASE, and CZ gates can be written as

$$\begin{aligned} H_j &= e^{i\pi/2} e^{-i\frac{\pi}{2} Z_j} e^{i\frac{\pi}{4} Y_j}, \\ S_j &= e^{i\pi/4} e^{-i\frac{\pi}{4} Z_j}, \\ CZ_{jk} &= e^{i\pi/4} e^{i\frac{\pi}{4} Z_j Z_k} e^{-i\frac{\pi}{4} Z_j} e^{-i\frac{\pi}{4} Z_k}. \end{aligned} \quad (8)$$

Similar identities can be derived for any Clifford gate acting on a fixed number of qubits.

---

**Algorithm 1:** Monte Carlo algorithm for improved randomized sparsification

---

**Input :** The CH-form  $(\omega_0, U_C^0, U_H^0)$  of the stabilizer state  $|\phi_y\rangle$ . A decomposition of the form in eq. (10) for  $i = 1, \dots, m$ .

**Output:** A random vector  $|z\rangle$  such that  $\mathbb{E}(|z\rangle) = |\psi\rangle$ .

```

1 begin
2   Set  $(\omega, U_C, U_H) \leftarrow (\omega_0, U_C^0, U_H^0)$ 
3   for  $i=1, \dots, m$  do
4     Choose  $V'_{ij}$  with probability  $p_i(j) \equiv |b_{ij}|/\|b_i\|_1$ ;
5     Update  $(\omega, U_C, U_H)$  under  $V'_{ij}$ ;
6     Set  $\omega \leftarrow (b_{ij}/|b_{ij}|)\omega$ ;
7   Output  $|z\rangle = \|b\|_1 w U_C U_H |s\rangle$ .
8 end

```

---

where

$$V'_{ij} \equiv \prod_k \exp(i\theta'_{ijk} P'_{ijk}), \quad (11)$$

$\theta'_{ijk} \in \{\pm\theta_{ijk}\}$ , and the  $P'_{ijk}$  are Pauli matrices each of which can be computed efficiently in time  $O(c)$ . Storing this description of  $U'$  requires  $O(nm)$  bits of memory, since each Pauli matrix requires  $O(n)$  bits to specify.

To conclude, we can write the output state of the circuit  $U$  as

$$\begin{aligned} |\psi\rangle &= U'|\phi\rangle \\ &= \sum_i a_i |u_i\rangle, \end{aligned} \quad (12)$$

where each stabilizer state  $|u_i\rangle$  is produced by acting on the fixed stabilizer state  $|\phi\rangle$  with a circuit consisting of  $O(m)$  Clifford gates of the form  $\exp(i\theta P)$ ,  $\theta \in (\pi/4)\mathbb{Z}$ . Combining this with Proposition 1 shows that we can produce the CH-form of  $|u_i\rangle$  in time  $O(mn^2)$ , which improves on the  $O(\ell n^2)$  required by the method described in [5].

This improvement is applicable to both the norm estimation and the heuristic Metropolis algorithms of [5]. The total time cost of computing the CH-forms of the approximate state is  $O(kmn^2)$  in our method, vs.  $O(k\ell n^2)$  in [5]. The runtime reduction factor of  $\ell/m$  is significant because the randomized sparsification subroutine requires computing an exponential number of CH-forms, roughly  $k \approx \|a\|_1^2 \epsilon^{-2}$ , in order to obtain an  $\epsilon$ -approximation to  $|\psi\rangle$ . Recall that the factor  $\|a\|_1^2$  is exponential in the number of non-Clifford gates. For a Clifford+T circuit, for example, it is equal to  $2^{0.228m}$ , where  $m$  is the T count.

**Monte-Carlo method.** Randomized sparsification (both the original version and our improved variant)

---

**Algorithm 2:** Adaptive Monte Carlo algorithm for randomized sparsification

---

**Input :** The CH-form  $(\omega_0, U_C^0, U_H^0)$  of the stabilizer state  $|\phi_y\rangle$ . A decomposition of the form in eq. (10) for  $i = 1, \dots, m$ .

**Output:** A random vector  $|z\rangle$  such that  $\mathbb{E}(|z\rangle) = |\psi\rangle$ .

```

1 begin
2   Set  $(\omega, U_C, U_H) \leftarrow (\omega_0, U_C^0, U_H^0)$ 
3   for  $i=1, \dots, m$  do
4     Compute eq. (15);
5     Choose  $l$  with probability  $p'_i(l) \equiv |a'_{il}|/\|a'_i\|_1$ ;
6     Set  $(\omega, U_C, U_H)$  to be the CH-form of  $|u_{il}\rangle$ ;
7     Set  $\omega \leftarrow (a'_{il}/|a'_{il}|)\omega$ ;
8   Output  $|z\rangle = \|b\|_1 w U_C U_H |s\rangle$ .
9 end

```

---

can be implemented as a Monte Carlo algorithm. The state space is the set of CH-tuples  $(\omega, U_C, U_H)$ . In our improved variant, the initial state is the CH-form of  $|\phi\rangle$ . For each gate  $U'_i$  in the reduced circuit, a random Clifford gate is applied to the current state according to eq. (10). Namely,  $V'_{ij}$  is applied with probability  $|b_{ij}|/\|b_i\|_1$ . The relative phases between the terms are accounted for by updating  $\omega$  appropriately. After the last gate, the final state is scaled by the factor  $\|b\|_1 \equiv \prod_i \|b_i\|_1$ . This is summarized in Algorithm 1. A direct calculation shows that the output  $|z\rangle$  of this algorithm satisfies

$$\mathbb{E}(|z\rangle) = |\psi\rangle. \quad (13)$$

Moreover, it is shown in [5] that the average of  $k$  runs of the algorithm is a random vector  $|\tilde{\psi}\rangle = k^{-1} \sum_{i=1}^k |z_i\rangle$  which satisfies

$$\Pr [\| |\tilde{\psi}\rangle - |\psi\rangle \| \geq \epsilon] \leq \frac{\|a\|_1^2/k}{\epsilon^2}. \quad (14)$$

For example, taking  $k = (2\|a\|_1/\epsilon)^2$  ensures that the probability of the error exceeding  $\epsilon$  is at most  $1/4$ .

**Adaptive Monte-Carlo.** The convergence of the estimator in the Monte Carlo of Algorithm 1 can be improved by adapting the probabilistic Clifford update to the particular state at each time step. Let the stabilizer state at the  $i$ th time step of the Monte Carlo be  $|u\rangle$ , and let

$$\begin{aligned} U'_i |u\rangle &= \sum_j a_{ij} V'_{ij} |u\rangle, \\ &= \sum_j a_{ij} |u_{ij}\rangle, \\ &= \sum_l a'_{il} |u_{il}\rangle, \end{aligned} \quad (15)$$

where all collinear terms in the second sum are collected in the third sum. Note that computing eq. (15) can be done efficiently by updating the CH description of each term under the Clifford gates  $V'_{ij}$ . It is then straightforward to show that  $\|a'_i\|_1 \leq \|a_i\|_1$ .

Consider the modified Monte Carlo in Algorithm 2. As before, we have  $\mathbb{E}(|z\rangle) = |\psi\rangle$ , and one can take the average of  $k$  iterations of this algorithm to approximate  $|\psi\rangle$ . As  $\|a'\|_1$  depends on the particular trajectory taken, we need to choose  $k \approx \Delta^2 \epsilon^{-2}$ , where

$$\Delta = \max_{\text{trajectories}} \|a'\|_1 \leq \|a\|_1. \quad (16)$$

Computing  $\Delta$  is generally hard since it involves a maximization over all trajectories. Unless a smaller upper bound for  $\Delta$  can be learned, the safe solution is to settle for the pessimistic sample size which scales like  $\|a\|_1^2$ . Even with this limitation, the above method improves the precision of the approximation obtained for the same sample size. We can use heuristics to predict whether adaptive sampling is useful. Let  $\alpha$  be the random variable which takes value  $\|a'\|_1$  with probability equal to the sum of the probabilities of all trajectories consistent with  $\|a'\|_1$ . One can estimate the mean and variance of  $\alpha$  by sampling, and heuristically argue that most trajectories have  $\|a'\|_1$  within, say, 3 standard deviations of the mean. Thus if the mean is much smaller than  $\|a\|_1$ , and the variance is not too large, it is a good indication that adaptive Monte Carlo improves the rate at which  $|\hat{\psi}\rangle$  converges to  $|\psi\rangle$ .

**Clifford decompositions of products of gates.** Finding a Clifford decomposition of  $U_i$  with minimal 1-norm  $\|a_i\|_1$  can be done using a convex optimization software [5, 12]. These brute-force optimizations can only be performed for one- or two-qubit gates, as it is too computationally intensive for more qubits.

For an operator  $A$ , denote by  $\Omega(A)$  the minimum 1-norm over all Clifford decompositions of  $A$ . Then  $\Omega$  is sub-multiplicative, i.e.

$$\Omega(AB) \leq \Omega(A)\Omega(B). \quad (17)$$

For two operators  $A$  and  $B$ , let us call a Clifford decomposition of the product  $AB$  *contractive* if its 1-norm is strictly less than  $\Omega(A)\Omega(B)$ . Finding a contractive Clifford decomposition of  $AB$  is generally hard unless  $A$  and  $B$  are at most 2-qubit gates acting on the same pair of qubits, so that we can use brute-force convex optimization. An exception is when the two operators can be mapped to the same pair of qubits via a Clifford  $C$ . In that case, one can find a contractive decomposition of  $CABC^\dagger$  using convex-optimization, and then apply the inverse  $C^\dagger$ , using the idea of eq. (10) to keep the terms in a form enabling fast CH updates. This method is well-suited to the case of Clifford+ $e^{i\theta P}$  circuits, which we discuss in the next section.

### 3 Example: Clifford+ $e^{i\theta P}$ circuits

As an application of the methods presented so far, we consider circuits where each non-Clifford gate has the form  $e^{i\theta P}$ , where  $P$  is a Pauli operator, and  $\theta \in [-\pi, \pi]$ . This class of circuits includes important cases such as Clifford+ $T$  circuits, and more generally Clifford+ $Z$  rotation circuits. For a circuit  $U$  of this type containing  $m$  non-Clifford gates, eq. (6) reduces to

$$U|0^n\rangle = e^{i\theta_m Q_m} \dots e^{i\theta_1 Q_1} |\phi\rangle, \quad (18)$$

for a list of Hermitian Pauli operators  $Q_1, \dots, Q_m$ , a list of angles  $\theta_1, \dots, \theta_m$ , and a stabilizer state  $|\phi\rangle$ . Note that the unitary  $e^{i\theta P}$  is Clifford if and only if  $\theta \in (\pi/4)\mathbb{Z}$ . Thus we may, without loss of generality, assume that  $\theta_i \in (0, \frac{\pi}{4})$ , as otherwise we can “factor out” Clifford gates  $e^{\pm i\frac{\pi}{4} Q_i}$  and absorb them into the state  $|\phi\rangle$ , at the cost of applying a Clifford mapping to the list of Pauli operators  $Q_j$  with  $j < i$ . The Clifford mapping can be quickly computed using the relation

$$e^{i\frac{\pi}{4} P} Q e^{-i\frac{\pi}{4} P} = \begin{cases} Q & \text{if } [P, Q] = 0 \\ iPQ & \text{if } \{P, Q\} = 0. \end{cases} \quad (19)$$

We refer to eq. (18) as the *canonical form* of the output state of a Clifford+ $e^{i\theta P}$  circuit.

For a single qubit Pauli operator  $P$ , the optimal decomposition of  $e^{i\theta P}$  into Clifford gates is given by [5]

$$e^{\pm i\theta P} = (\cos \theta - \sin \theta)\mathbb{1} + (\sqrt{2} \sin \theta) e^{\pm i\frac{\pi}{4} P}. \quad (20)$$


This decomposition achieves the minimum 1-norm,

$$\Omega(e^{i\theta P}) = \cos \theta + (\sqrt{2} - 1) \sin \theta \quad (0 \leq \theta \leq \frac{\pi}{4}). \quad (21)$$


This holds for  $n$ -qubit Pauli operators as well, which follows from i) invariance of  $\Omega$  under multiplication by Clifford unitaries, ii) invariance of  $\Omega$  under tensoring by the identity, i.e.  $\Omega(U \otimes \mathbb{1}) = \Omega(U)$ .

**Clifford decompositions of products of exponentiated Pauli operators.** If a sequence of exponentials in eq. (18) multiplies to a Clifford gate, it can be commuted past the gates prior to it and absorbed into the input state, effectively reducing the number of non-Clifford gates in the circuit. As we are dealing with a universal set of gates, it is generally intractable to determine whether such a product belongs to the Clifford group. In the simplest case of multiplying two gates  $e^{i\theta P} e^{i\phi Q}$  with  $\theta, \phi \in (0, \pi/4)$ , it can be shown that the product is a Clifford gate if and only if  $PQ = \pm \mathbb{1}$  and  $\theta + \phi \in (\pi/4)\mathbb{Z}$ . If  $PQ \neq \pm \mathbb{1}$ , there exists no better decomposition than the trivial one, obtained by decomposing each gate separately;

$$\Omega(e^{i\theta_1 P} e^{i\theta_2 Q}) = \Omega(e^{i\theta_1 P}) \Omega(e^{i\theta_2 Q}). \quad (22)$$



(a) One qubit



(b) Two qubits

Figure 1: Theoretical time cost of simulating various one- and two-qubit exponential sequences. The angles are chosen to be equal for simplicity. For a sequence of length  $L$ , the function  $f(\theta)$  is the 1-norm of the decomposition of the sequence into Clifford gates, and  $f(\theta)^{1/L}$  is the effective 1-norm per gate. The function  $2 \log(f(\theta)^{1/L})$  is the contribution of each gate to  $\alpha$ , where  $2^\alpha$  is the exponential scaling factor in the runtime of the simulation. The data is obtained using a convex optimization software (CVX with the SDPT3 solver).

This can be verified numerically for 1 and 2 qubits, and the general case follows from the fact that any pair of  $n$ -qubit Paulis  $P$  and  $Q$  can be mapped by a Clifford circuit to a pair of Paulis with support on the first two qubits only, together with the fact that  $\Omega(U \otimes \mathbb{1}) = \Omega(U)$ .

For products of three or more gates, we can find decompositions with significantly reduced 1-norm in certain cases. Numerically, we observe that the possibility of 1-norm reduction is correlated with the rank of the binary representation of the Paulis. Specifically, for  $x, z \in \mathbb{Z}_2^n$  write  $X[x] \equiv X^{x_1} \otimes \dots \otimes X^{x_n}$ ,  $Z[z] \equiv Z^{z_1} \otimes \dots \otimes Z^{z_n}$ , and  $P_{(x,z)} \equiv i^{x \cdot z} X[x]Z[z]$ . Then numerics in the two qubit case suggest that,

$$\Omega \left( \prod_j e^{i\theta_j P_{v_j}} \right) < \prod_j \Omega(e^{i\theta_j P_{v_j}}) \quad (23)$$

if and only if  $\text{rank}(v_1, v_2, \dots, v_k) < k$  with  $\text{rank}(\cdot)$  computed mod 2. For example, it is possible to numerically

verify that

$$\Omega(e^{i\theta_1 X} e^{i\theta_2 Y} e^{i\theta_3 Z}) < \Omega(e^{i\theta_1 X}) \Omega(e^{i\theta_2 Y}) \Omega(e^{i\theta_3 Z}), \quad (24)$$

for any choice of angles  $(\theta_1, \theta_2, \theta_3)$ . We can also find the optimal Clifford expansion<sup>2</sup>

$$e^{i\theta_1 X} e^{i\theta_2 Y} e^{i\theta_3 Z} = \sum_i \alpha_i C_i. \quad (26)$$

Therefore for any triple of  $n$ -qubit Paulis  $(P, Q, W)$  which are equivalent to  $(X_1, Y_1, Z_1)$  via a Clifford  $V$ , the optimal decomposition is given by

$$e^{i\theta_1 P} e^{i\theta_2 Q} e^{i\theta_3 W} = \sum_i \alpha_i V^\dagger (C_i \otimes \mathbb{1}^{\otimes n-1}) V. \quad (27)$$

As before, each Clifford term  $V^\dagger (C_i \otimes \mathbb{1}^{\otimes n-1}) V$  is easy to cast as a product of a small number of Cliffords of the form  $e^{i\phi R}$  where  $R$  is a Pauli and  $\phi \in (\pi/4)\mathbb{Z}$ . This way we obtain a decomposition of the many-body operator  $e^{i\theta_1 P} e^{i\theta_2 Q} e^{i\theta_3 W}$  which is both optimal and in a form enabling fast phase-sensitive updates.

The reduction in the 1-norm by decomposing products of gates can be quite large, see for example Figure 1, where the log of the 1-norm is plotted for different one and two qubit gate sequences.

## 4 Amplitude estimation

To illustrate the value of Clifford recompilation, we consider the problem of estimating circuit amplitudes of near-Clifford circuits. Randomized sparsification can be used to estimate the amplitudes  $\langle x|\psi \rangle$  of the output state. Indeed eq. (13) implies that

$$\mathbb{E}(\langle x|z \rangle) = \langle x|\psi \rangle, \quad (28)$$

where  $|z \rangle$  is the random vector generated by Algorithm 1. By a complex-valued version of Hoeffding's inequality [11, 16], the random vector  $|\tilde{\psi} \rangle$  obtained by averaging  $k$  instances of  $|z \rangle$  satisfies

$$\Pr [|\langle x|\psi \rangle - \langle x|\tilde{\psi} \rangle| \geq \epsilon] \leq \delta, \quad (29)$$

provided we take  $k = O(\|a\|_1^2 \epsilon^{-2} \log(\delta^{-1}))$ . Thus by averaging enough evaluations of  $\langle x|z \rangle$ , one obtains a

<sup>2</sup>For  $0 \leq \theta \leq \pi/8$ , the expansion

$$e^{i\theta X} e^{i\theta Y} e^{i\theta Z} = (\cos 3\theta - \sin \theta) \mathbb{1} + e^{i\pi/4} [\sin 2\theta (\cos \theta - \sin \theta)] SH + (\sqrt{2} \sin 2\theta \sin \theta) HX + (\sqrt{2} \sin 2\theta \sin \theta) SHS. \quad (25)$$

is optimal. More generally, the Clifford terms in eq. (26) vary depending on  $(\theta_1, \theta_2, \theta_3)$ , and a general expression is too tedious to write down, although finding the optimal decomposition numerically is easy on a case-by-case basis.

good estimate of  $\langle x|\psi\rangle$  with high probability. Given the CH-form of  $|z\rangle$ , it takes time  $O(n^2)$  to compute the amplitude  $\langle x|z\rangle$  using the methods in [5]. This gives a runtime scaling of  $O(\|a\|_1^2 mn^2 \epsilon^{-2} \log(\delta^{-1}))$  for computing an additive error estimate of the amplitude, using the improved sparsification Monte Carlo of Section 2. Typically we are interested in estimating the probability  $\Pr(x) = |\langle x|\psi\rangle|^2$ , rather than the amplitude. It is straightforward to check that if  $|\langle x|\tilde{\psi}\rangle - \langle x|\psi\rangle| \leq c\epsilon$ , where  $c = \sqrt{2} - 1$ , then  $||\langle x|\tilde{\psi}\rangle|^2 - P(x)| \leq \epsilon$  (as long as  $\epsilon \leq 1$ ). In other words, estimating the probability instead of the amplitude only requires multiplying  $k$  by the constant  $c^{-2} \approx 6$ .

Note that the heuristic Metropolis simulator of [5] relies on taking the ratio between two amplitude estimates, obtained exactly as described above, in order to determine the next move at each Metropolis step.

Estimating amplitudes as above relies on very few properties of stabilizer states. It is only needed that stabilizer states admit an efficient, phase-sensitive description which can be updated under Clifford gates, and that this description can be used to efficiently compute any desired amplitude of the state. Thus estimating amplitudes as described in this section, and indeed the heuristic Metropolis simulator of [5], can both be used for other states and gates satisfying similar criteria. Perhaps the main candidate for this generalization is Fermionic Gaussian states and nearest-neighbor matchgates [17, 18]. Nearest-neighbor matchgates acting on Fermionic Gaussian states are known to be classically simulable, and yet augmenting the gate set with nearest-neighbor SWAP gates is sufficient for universal quantum computing. This motivates searching for matchgate-decompositions of the SWAP gate. Our numerical search suggests that the decomposition

$$\text{SWAP} = \frac{e^{i\pi/4}}{\sqrt{2}}G(\mathbb{1}, -iX) + \frac{e^{-i\pi/4}}{\sqrt{2}}G(\mathbb{1}, iX) \quad (30)$$

has minimal 1-norm over all matchgate-decompositions of SWAP, where  $G(\mathbb{1}, \pm iX)$  is the  $\pm i$ SWAP gate (which is a matchgate). This decomposition has 1-norm equal to  $\sqrt{2}$ , which implies that amplitude estimation using this decomposition has a runtime scaling like  $O(2^m)$ , where  $m$  is the number of SWAP gates. Since it is possible to exactly compute the amplitude in time  $O(2^m)$  by computing the  $2^m$  matchgate contributions corresponding to the expansion in eq. (30), this decomposition is not useful for this purpose. It may be possible nevertheless to find smaller 1-norms by considering decompositions of parallel SWAP gates into nearest-neighbor matchgate circuits involving more than two qubits.

Applying the same idea to local vs. entangling gates,

we find numeric evidence that the decomposition

$$\text{CPHASE} = \frac{e^{i\pi/4}}{\sqrt{2}}(S_1^\dagger S_2^\dagger - iS_1 S_2) \quad (31)$$

achieves the minimum 1-norm over all decompositions of CPHASE into product gates (not necessarily Cliffords), where  $S_i$  is the PHASE gate on qubit  $i$ . Similar to the matchgate case, the 1-norm of this decomposition is too large to provide an improvement over brute force methods. Nevertheless, our numerics is not exhaustive and a better decomposition may exist. It may also be the case that a decomposition of the form

$$\text{CPHASE}^{\otimes k} = \sum_j a_j V_j, \quad (32)$$

exists such that each  $V_j$  is a product gate on  $2k$  qubits and  $\|a\|_1 < 2^{k/2}$ . Our numerical search did not detect such decompositions for  $k = 2$  allowing for up to 4 terms in the sum.

## 5 Conclusion

We have described how to improve the simulators of Bravyi et al [5] by recompiling the circuit using exponentiated Pauli operators. The theoretical runtime reduction is given by the factor  $\ell/m$  where  $\ell$  and  $m$  are respectively the total and non-Clifford gate counts, which can be significant when the target circuit is vastly dominated by Clifford gates. As an added benefit, the recompilation puts the circuit in a form where it is easier to search for optimal Clifford decompositions, since all non-Clifford gates appear in sequence. This can significantly reduce the exponential prefactor in the runtime. The circuit recompilation is well-suited to be used with a scheme for computing additive-error estimates of outcome probabilities. This scheme requires much less work than the full techniques of [5], and can in principle be used with a broader class of states and gates.

**Acknowledgements.** This research was supported by the U.S. Army Research Office through grant W911NF-14-1-0103, and by funding from the Government of Ontario through TQT, and the Government of Canada through CFREF and NSERC.

## References

- [1] Ryan S Bennink, Erik M Ferragut, Travis S Humble, Jason A Laska, James J Nutaro, Mark G Pleszkoch, and Raphael C Pooser. Monte carlo simulation of near-clifford quantum circuits. *arXiv preprint*, 2017. URL <https://arxiv.org/abs/1703.00111>.

- [2] Sergio Boixo, Sergei V Isakov, Vadim N Smelyanskiy, and Hartmut Neven. Simulation of low-depth quantum circuits as complex undirected graphical models. *arXiv preprint arXiv:1712.05384*, 2017. URL <https://arxiv.org/abs/1712.05384>.
- [3] Sergey Bravyi and David Gosset. Improved classical simulation of quantum circuits dominated by clifford gates. *Phys. Rev. Lett.*, 116:250501, Jun 2016. DOI: [10.1103/PhysRevLett.116.250501](https://doi.org/10.1103/PhysRevLett.116.250501).
- [4] Sergey Bravyi, Graeme Smith, and John A Smolin. Trading classical and quantum computational resources. *Physical Review X*, 6(2):021043, 2016. DOI: [10.1103/physrevx.6.021043](https://doi.org/10.1103/physrevx.6.021043).
- [5] Sergey Bravyi, Dan Browne, Padraic Calpin, Earl Campbell, David Gosset, and Mark Howard. Simulation of quantum circuits by low-rank stabilizer decompositions. *arXiv preprint*, 2018. URL <https://arxiv.org/abs/1808.00128>.
- [6] Jianxin Chen, Fang Zhang, Mingcheng Chen, Cupjin Huang, Michael Newman, and Yaoyun Shi. Classical simulation of intermediate-size quantum circuits. *arXiv preprint*, 2018. URL <https://arxiv.org/abs/1805.01450>.
- [7] Zhaoyun Chen, Qi Zhou, Cheng Xue, Xia Yang, Guangcan Guo, and Guoping Guo. 64-qubit quantum circuit simulation. *Science Bulletin*, 63(15):964–971, aug 2018. DOI: [10.1016/j.scib.2018.06.007](https://doi.org/10.1016/j.scib.2018.06.007).
- [8] Hans De Raedt, Fengping Jin, Dennis Willsch, Madita Nocon, Naoki Yoshioka, Nobuyasu Ito, Shengjun Yuan, and Kristel Michielsen. Massively parallel quantum computer simulator, eleven years later. *arXiv preprint arXiv:1805.04708*, 2018. DOI: [10.1016/j.cpc.2018.11.005](https://doi.org/10.1016/j.cpc.2018.11.005).
- [9] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Physical Review A*, 86(3), sep 2012. DOI: [10.1103/physreva.86.032324](https://doi.org/10.1103/physreva.86.032324).
- [10] Aram W. Harrow and Ashley Montanaro. Quantum computational supremacy. *Nature*, 549(7671): 203–209, sep 2017. DOI: [10.1038/nature23458](https://doi.org/10.1038/nature23458).
- [11] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association*, 58(301):13–30, 1963. DOI: [10.1007/978-1-4612-0865-5\\_26](https://doi.org/10.1007/978-1-4612-0865-5_26).
- [12] Mark Howard and Earl Campbell. Application of a resource theory for magic states to fault-tolerant quantum computing. *Physical Review Letters*, 118(9):090501, 2017. DOI: [10.1103/physrevlett.118.090501](https://doi.org/10.1103/physrevlett.118.090501).
- [13] C. Neill, P. Roushan, K. Kechedzhi, S. Boixo, S. V. Isakov, V. Smelyanskiy, A. Megrant, B. Chiaro, A. Dunsworth, K. Arya, R. Barends, B. Burkett, Y. Chen, Z. Chen, A. Fowler, B. Foxen, M. Giustina, R. Graff, E. Jeffrey, T. Huang, J. Kelly, P. Klimov, E. Lucero, J. Mutus, M. Neeley, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, H. Neven, and J. M. Martinis. A blueprint for demonstrating quantum supremacy with superconducting qubits. *Science*, 360(6385): 195–199, apr 2018. DOI: [10.1126/science.aao4309](https://doi.org/10.1126/science.aao4309).
- [14] John Preskill. Quantum computing in the NISQ era and beyond. *Quantum*, 2:79, aug 2018. DOI: [10.22331/q-2018-08-06-79](https://doi.org/10.22331/q-2018-08-06-79).
- [15] Mikhail Smelyanskiy, Nicolas PD Sawaya, and Alán Aspuru-Guzik. qhipster: the quantum high performance software testing environment. *arXiv preprint arXiv:1601.07195*, 2016.
- [16] Dan Stahlke. Quantum interference as a resource for quantum speedup. *Phys. Rev. A*, 90:022302, Aug 2014. DOI: [10.1103/PhysRevA.90.022302](https://doi.org/10.1103/PhysRevA.90.022302).
- [17] Barbara M. Terhal and David P. DiVincenzo. Classical simulation of noninteracting-fermion quantum circuits. *Physical Review A*, 65(3), mar 2002. DOI: [10.1103/physreva.65.032325](https://doi.org/10.1103/physreva.65.032325).
- [18] Leslie G Valiant. Quantum circuits that can be simulated classically in polynomial time. *SIAM Journal on Computing*, 31(4):1229–1254, 2002. DOI: [10.1137/s0097539700377025](https://doi.org/10.1137/s0097539700377025).