Efficient decomposition of single-qubit gates into V basis circuits

Alex Bocharov,[1,*] Yuri Gurevich,[2,†] and Krysta M. Svore[1,‡]

[1]*Quantum Architectures and Computation Group, Microsoft Research, Redmond, Washington 98052, USA*

[2]*Research in Software Engineering Group, Microsoft Research, Redmond, Washington 98052, USA*

(Received 28 March 2013; published 12 July 2013)

We develop efficient algorithms for compiling single-qubit unitary gates into circuits over the universal V *basis*. The V *basis* is an alternative universal basis to the more commonly studied basis consisting of Hadamard and $\pi/8$ gates. We propose two classical algorithms for quantum circuit compilation: the first algorithm has expected polynomial time [in precision $\log(1/\epsilon)$] and produces an ϵ approximation to a single-qubit unitary with a circuit depth $\leqslant 12\log_5(2/\epsilon)$. The second algorithm performs optimized direct search and yields circuits a factor of 3 to 4 times shorter than our first algorithm, but requires time exponential in $\log(1/\epsilon)$; however, we show that in practice the runtime is reasonable for an important range of target precisions. Decomposing into the V basis may offer advantages when considering the fault-tolerant implementation of quantum circuits.

DOI: 10.1103/PhysRevA.88.012313 PACS number(s): 03.67.Lx, 03.65.Fd

I. INTRODUCTION

Determining the optimal fault-tolerant compilation, or decomposition, of a quantum gate is critical for designing a quantum computer. Decomposition of single-qubit unitary gates into the $\{H,T\}$ basis, where H is the Hadamard gate and T is the rotation about the z axis by $\pi/4$, has been well studied in recent years. However, there have been few studies of decomposing into alternative bases, which may offer significant improvements in circuit depth or resource cost. We consider the task of decomposing a single-qubit unitary gate into a sequence of gates drawn from the V *basis*, first introduced in [1,2]. The V basis contains rotations by $\theta = \cos^{-1}(-3/5)$ around the x, y, and z axis. For example, the matrix for the z-axis rotation, denoted as V_3, is given by

$$V_3 = \frac{1}{\sqrt{5}}\begin{bmatrix} 1+2i & 0 \\ 0 & 1-2i \end{bmatrix} = \frac{1+2i}{\sqrt{5}}\begin{bmatrix} 1 & 0 \\ 0 & \frac{-(3+4i)}{5} \end{bmatrix}.$$

The x-axis rotation $V_1 = HV_3H$ and the y-axis rotation $V_2 = SV_1S^\dagger$ are both Clifford adjoints of V_3, where S is the rotation around the z axis by π. Historically, the V basis was the first basis proven to be *efficiently universal*, meaning that the decomposition sequence is guaranteed to be of depth $O(\log(1/\epsilon))$ [3], however the proof did not offer a constructive method for finding such sequences. In this work, we provide two decomposition algorithms that yield short circuits for the universal V basis.

Recently, it has been shown that $\{H,T\}$ is also efficiently universal [4,5], and the proofs are constructive. Notably, characterization of $\{H,T\}$ circuits [6,7] has led to a constructive algorithm for an efficient ancilla-free compilation to precision ϵ of a given single-qubit unitary into the $\{H,T\}$ basis with a maximum T-count guarantee of $4\log_2(1/\epsilon)+11$ for Z rotations and $12\log_2(1/\epsilon)+K$, where $K \sim 33$ for general unitaries [4]. Further improvements were given in Ref. [5] in the form of a decomposition algorithm that outputs shorter ancilla-free approximation circuits (albeit with less efficient runtime) with an expected T count of $9.63\log_2(1/\epsilon) - 20.79$. These results are plotted in Fig. 1(a), for T count versus precision ϵ. The solid blue curve plots the theoretical bound for the algorithm given in Ref. [4]. The dashed red curve is based on interpolation of the experimental results given in Ref. [5]. These algorithms seek to minimize the number of T gates due to the high fault-tolerant implementation cost of T, as compared to the relatively inexpensive cost of Clifford gates.

Given that the V basis is efficiently universal [3], it is natural to consider constructing an efficient decomposition algorithm for it. We present two algorithms for compiling efficiently into the V basis, where we seek to minimize the number of non-Clifford V gates due to their high implementation costs. The first algorithm approximates single-qubit unitaries over the set consisting of the V basis and the Clifford group; the second approximates over the set consisting of the V basis and the Pauli gates. The first algorithm (Sec. III) is a randomized algorithm that runs in expected polynomial time and delivers ϵ approximations with V count $\leqslant 12\log_5(2/\epsilon)$. The second algorithm (Sec. IV) is based on direct search and produces ϵ approximations with V count $\leqslant 3\log_5(1/\epsilon)$ for most single-qubit unitaries, and approximations with V count $4\log_5(2/\epsilon)$ for edge cases. The compilation time is linear in $1/\epsilon$ and thus exponential in $\log(1/\epsilon)$, however, in practice we find extremely short circuits (of length $L = 28$) at precision level $\epsilon = 3\times 10^{-7}$ with merely 1 min of classical CPU time and modest space usage.

These results are plotted, for V count versus precision ϵ, in Fig. 1(b), for our randomized algorithm (solid red line), an empirical version of our randomized algorithm (dashed green line), and our direct search algorithm (double blue curve); the curves are averages over the decomposition of 1000 random unitaries. The results indicate that efficient approximating circuits can be obtained when considering other universal bases, and that there may be potential advantages in decomposing into other such universal bases. Our work gives an alternative to $\{H,T\}$ decomposition and produces

*alexeib@microsoft.com
†gurevich@microsoft.com
‡ksvore@microsoft.com

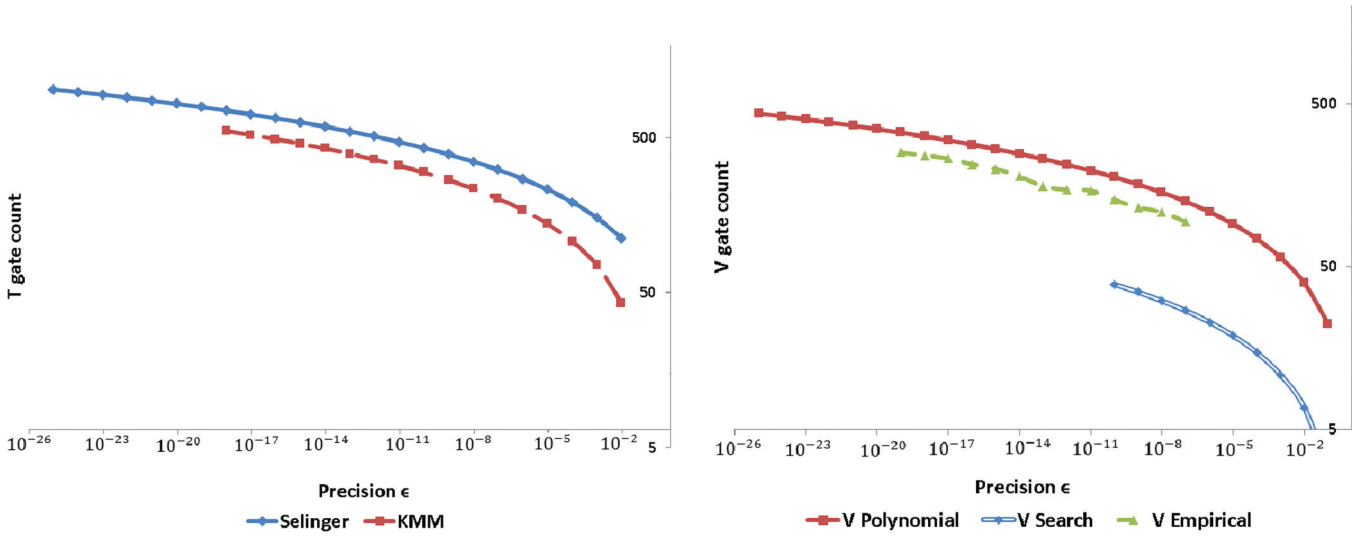

FIG. 1. (Color online) (a) T count versus precision ϵ for decomposing using the algorithm in Ref. [4] (solid upper curve) and the algorithm in Ref. [5] (dashed lower curve). (b) V count versus precision ϵ for decomposition into the V basis using the randomized algorithm (Sec. III; solid upper curve), an empirical tightening of that algorithm (dashed middle curve), and the direct search algorithm (Sec. IV; double lower curve).

circuits with lengths matching the proven lower bound of $\Omega(\log(1/\epsilon))$ [3].

Our motivation for studying the V basis is twofold. First, we have been motivated by the efficient universality proof in [3] and by the challenge of finding an algorithmic version of it. Second, there have been advances in fault-tolerant implementations of the V gates. A recent protocol for distillation of nonstabilizer states [8] allows a V gate to be implemented by way of T gates. Another approach is to use the circuit shown in Fig. 2 (slightly modified from Ex. 4.41 in [9]) in conjunction with state-of-the-art protocols for distilling Toffoli states [10–12]. These implementation methods are presented in Appendix.

In the case of both the $\{H,T\}$ basis and the V basis, the non-Clifford gates cannot be implemented fault tolerantly in a purely unitary fashion. State-of-the-art fault-tolerant implementations of T are based on state distillation, which requires significant use of measurement, classical feedback, and ancillae qubits [13–15]. Similarly, existing fault-tolerant implementations of the V gates (see Appendix) also require measurement, classical feedback, and ancillae qubits. In both cases, efficient fault-tolerant methods continue to be developed and optimized. For this reason, we choose to cost the decomposition algorithms in terms of the number of non-Clifford gates. A fine-grained cost comparison can then be determined based on a chosen fault-tolerant implementation or device architecture.

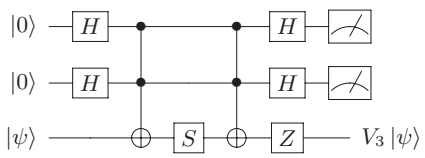

FIG. 2. Circuit to perform $V_3|\psi\rangle$ when both measurement outcomes are 0, or $|\psi\rangle$ otherwise [9].

The paper is organized as follows. Section II outlines preliminary definitions and theorems. Section III presents our randomized algorithm and corresponding experimental results. We also compare number theory techniques used at the core of our algorithm with those found in [4,5]. Our direct search algorithm is presented in Sec. IV. Conclusions and future directions are discussed in Sec. V.

II. PRELIMINARIES

The efficiently universal single-qubit unitary basis introduced in Refs. [1] and [2] and further developed in Ref. [3] consists of the following six special unitaries:

$$V_1 = (I + 2\,i\,X)/\sqrt{5}, \quad V_1^{-1} = (I - 2\,i\,X)/\sqrt{5},$$
$$V_2 = (I + 2\,i\,Y)/\sqrt{5}, \quad V_2^{-1} = (I - 2\,i\,Y)/\sqrt{5},$$
$$V_3 = (I + 2\,i\,Z)/\sqrt{5}, \quad V_3^{-1} = (I - 2\,i\,Z)/\sqrt{5}.$$

We call this basis the V *basis*.

The subgroup $\langle V\rangle \subset \mathrm{SU}(2)$ generated by this basis is everywhere dense in SU(2) and thus $\{V_i, V_i^{-1}, i = 1,2,3\}$ is a universal basis. Let the set of W *circuits* be the set of those circuits generated by this basis and the Pauli matrices I,X,Y,Z.

It is important to note that the monoid $\langle W\rangle = \langle X,Y,Z, V_1,V_2,V_3\rangle \subset \mathrm{SU}(2)$ contains all of the $\{V_i^{-1}, i = 1,2,3\}$ and thus is in fact a subgroup of SU(2) containing $\langle V\rangle$.

W circuits constitute a slight liberalization of the approach in Ref. [3], where only circuits in the V basis are considered. Our justification for the liberalization is that the Pauli operators are a staple of any quantum computing architecture and can be implemented fault tolerantly at a very low resource cost in comparison to a non-Clifford group gate. It is also noted that

the single-qubit Clifford group $\mathcal{C}$ in combination with *any* of the six V matrices generates a monoid $\langle \mathcal{C} + V \rangle \subset \mathrm{SU}(2)$ that is in fact a group, containing $\langle W \rangle$.

We call the number of V gates the V *count* of a circuit and denote it as V_c. It is easy to show that an irreducible W circuit contains at most one nonidentity Pauli gate. Thus, if v is the V count of such a circuit, then the overall depth of the circuit is either v or $v+1$.

Throughout, we use *trace distance* to measure the distance between two unitaries $U, V \in \mathrm{PSU}(2)$:

$$\mathrm{dist}(U,V) = \sqrt{1 - |\mathrm{tr}(U V^\dagger|/2}, \tag{1}$$

and denote the distance between a target unitary and the approximating unitary as *precision* ϵ.

According to [3], any single-qubit unitary can be approximated to a given precision ϵ by a V circuit with V count $O(\log(\frac{1}{\epsilon}))$, however the proof in Ref. [3] is nonconstructive, and no algorithm for actual synthesis of the approximating circuits has yet been shown. Here we develop effective solutions for synthesizing W-circuit approximations of single-qubit unitaries.

Our solutions are based on the following theorem:

Theorem 1. A single-qubit unitary gate U can be exactly represented as a W circuit of V count $V_c \leqslant L$ if and only if it has the form

$$U = (a\, I + b\, i\, X + c\, i\, Y + d\, i\, Z)\, 5^{-L/2}, \tag{2}$$

where a,b,c,d are integers such that $a^2 + b^2 + c^2 + d^2 = 5^L$.

Theorem 1 follows from Theorem 2 given below, which also gives rise to a constructive procedure for synthesizing a W circuit that represents such a U. We begin by sketching a linear-time subalgorithm for exact W-circuit synthesis that employs arithmetic of Lipschitz quaternions [16,17]. More specifically, consider the group W of quaternions generated by

$$\pm 1, \pm \mathbf{i}, \pm \mathbf{j}, \pm \mathbf{k}, 1 \pm 2\mathbf{i},\ 1 \pm 2\mathbf{j}, 1 \pm 2\mathbf{k}. \tag{3}$$

Then the following holds:

Theorem 2. (1) W is equal to the set of Lipschitz quaternions with norms $5^l, (l \in \mathbb{Z}, l \geqslant 0)$. (2) Consider the group $W_1 = \{w/\sqrt{\mathrm{norm}(w)} \mid w \in W\}$. Then the subgroup of gates in PSU(2) representable as exact W circuits is isomorphic to the central quotient $W_1/Z(W_1)$ where $Z(W_1) = \mathbb{Z}_2 = \{1,-1\}$.

Proof. (1) We recall that the quaternion norm is multiplicative and that $\pm 1, \pm\mathbf{i}, \pm\mathbf{j}, \pm\mathbf{k}$ are the only Lipschitz quaternions of norm 1. Thus statement (1) is true for $l = 0$.

We prove it for $l = 1$: More specifically, let $q = a + b\,\mathbf{i} + c\,\mathbf{j} + d\,\mathbf{k}, a,b,c,d \in \mathbb{Z}$ and $\mathrm{norm}(q) = a^2 + b^2 + c^2 + d^2 = 5$.

Decompositions of 5 into sums of squares of four integers are easily enumerated and we conclude that exactly two of the coefficients in the list $\{a,b,c,d\}$ are zero, exactly one is ± 1, and exactly one is ± 2.

If $a = \pm 1$ then we observe that q is equal to one of $1 \pm 2\mathbf{i}, 1 \pm 2\mathbf{j}, 1 \pm 2\mathbf{k}, -(1 \pm 2\mathbf{i}), -(1 \pm 2\mathbf{j}), -(1 \pm 2\mathbf{k})$ and thus belongs to W.

If one of b,c,d is ± 1 we reduce the proof to the previous observation by multiplying q times one of $\mathbf{i},\mathbf{j},\mathbf{k}$.

For example, if $c = \pm 1$, then the real part of $-\mathbf{j}q$ is equal to $c = \pm 1$.

Consider now a quaternion q with $\mathrm{norm}(q) = 5^l, l \geqslant 1$.

Let $q = p_1 \ldots p_m$ be a prime quaternion factorization of q. Since $5^l = \mathrm{norm}(q) = \mathrm{norm}(p_1) \ldots \mathrm{norm}(p_m)$, for each $i = 1, \ldots m$, the $\mathrm{norm}(p_i)$ is either 5 or 1. As we have shown above (considering $l = 0,1$), in either case $p_i \in W$.

(2) Effective homomorphism h of W_1 onto the W circuits is the multiplicative completion of the following map:

$$\begin{aligned}
\mathbf{i} &\to i\, X, \\
\mathbf{j} &\to i\, Y, \\
\mathbf{k} &\to i\, Z, \\
(1 \pm 2\mathbf{i})/\sqrt{5} &\to (1 \pm 2\, i\, X)/\sqrt{5}, \\
(1 \pm 2\mathbf{j})/\sqrt{5} &\to (1 \pm 2\, i\, Y)/\sqrt{5}, \\
(1 \pm 2\mathbf{k})/\sqrt{5} &\to (1 \pm 2\, i\, Z)/\sqrt{5}.
\end{aligned}$$

The correctness of this definition of homomorphism h is verified by direct comparison of multiplicative relations between the generators of W_1 and $g(W) = \{i\, X, i\, Y, i\, Z, (1 \pm 2\, i\, X)/\sqrt{5}, (1 \pm 2\, i\, Y)/\sqrt{5}, (1 \pm 2\, i\, Z)/\sqrt{5}\}$. These relations happen to be identical.

Effective homomorphism h is an epimorphism since all of the generators $g(W)$ of the W-circuits group are by design in its image.

The characterization of $\mathrm{Ker}(h)$ is derived from representation of quaternions as orthogonal rotations of the three-dimensional Euclidean space.

Let us arbitrarily map the units $\mathbf{i},\mathbf{j},\mathbf{k}$ into vectors of an orthonormal basis in the Euclidean space and let us label the corresponding basis vectors $e(\mathbf{i}), e(\mathbf{j}), e(\mathbf{k})$. For a quaternion with zero real part $p = b\,\mathbf{i} + c\,\mathbf{j} + d\,\mathbf{k}$ we write $e(p) = b * e(\mathbf{i}) + c * e(\mathbf{j}) + d * e(\mathbf{k})$.

Let H_1 be the group of quaternions of norm 1 and $g : H_1 \to \mathrm{SO}(3)$ be the representation defined as $g(q)[e(b)] = e(q * b * q^{-1})$.

It is known [17] that $g(q)$ is an orthogonal rotation; g is a representation of the group of quaternions of norm 1 and that the kernel of this representation is the cyclic group $\mathbb{Z}_2 = \{1,-1\}$.

The group of quantum gates PSU(2) also has a standard orthogonal representation stemming from its adjoint representation on the Lie algebra. More specifically if psu(2) is regarded as the algebra of zero-trace Hermitian matrices then $ad : \mathrm{PSU}(2) \to \mathrm{Aut}(\mathrm{psu}(2))$, where Aut is the automorphism group, is $ad(u)[m] = u\, m\, u^{-1}$.

The adjoint representation of PSU(2) is faithful.

If we regard the above homomorphism h as the homomorphism $h : W_1 \to \mathrm{PSU}(2)$ then it is immediate that $ad\, h = g$ on W_1. Since ad is faithful, i.e., injective, the kernel of h coincides with $\mathrm{Ker}(g) = Z(W_1) = (\mathbb{Z})_2 = \{-1,1\}$. ■

Lipschitz quaternions form a division ring, and in view of Theorem 2, a quaternion with norm equal to 5^l can be decomposed into a product of generators in Eq. (3) in l trial division steps. The *decomposition subalgorithm* (Algorithm 1) is thus as follows, where the input is a Lipschitz quaternion q of norm 5^l:

Algorithm 1 Decomposition subalgorithm

Require: A quaternion q with norm 5^l
ret $\leftarrow$ empty list
while norm$(q) > 0$ **do**
  find d in $\{1 \pm 2\mathbf{i}, 1 \pm 2\mathbf{j}, 1 \pm 2\mathbf{k}\}$ such that d divides q
  ret $\leftarrow \{d\} + ret$
  $q \leftarrow q/d$ //divides norm(q) by 5
end while
if $q \neq 1$ **then**
  ret $\leftarrow q +$ ret
end if
return ret

Now, given a unitary U as described in Theorem 1, we associate with it the quaternion $q = a + b\,\mathbf{i} + c\,\mathbf{j} + d\,\mathbf{k}$ that has norm 5^L and thus belongs to the subgroup W. It is easy to translate the factorization of q in the basis given in Eq. (3) into a factorization of U in the W basis. Thus, the approximation of a target unitary gate G by a W circuit is constructively reduced to approximating G with a unitary U as described in Theorem 1.

III. RANDOMIZED APPROXIMATION ALGORITHM

In this section, we present an algorithm for decomposing single-qubit unitaries into a circuit in the set $\langle \mathcal{C} + V \rangle$, where $\mathcal{C}$ is the set of single-qubit Clifford gates and V is one of the V gates. The expected polynomial runtime is based on a conjecture, for which we have developed empirical evidence based on computer simulation. We first present the conjecture and relevant number theory background, and then present the compilation algorithm.

A. Number theory background

Let N be a large positive integer, and Δ be a relatively small fixed offset value. Let x, y be standard coordinates on a two-dimensional Euclidean plane. We introduce the circumference

$$C(N,\Delta) = \{(x,y) \mid x^2 + y^2 = (\sqrt{N} - \Delta)^2\}.$$

Let $R(N,\Delta)$ be the circular ring of width Δ defined as

$$R(N,\Delta) = \{(x,y) \mid (\sqrt{N} - \Delta)^2 < x^2 + y^2 < N\}.$$

Consider a tangent straight line at any point on the circumference $C(N,\Delta)$. The line divides the plane into two half planes; let P_+ be the half plane that does not contain the origin. Next, define the circular segment

$$A(N,\Delta,P_+) = R(N,\Delta) \bigcap P_+.$$

The ring $R(N,\Delta)$ and the circular segment $A(N,\Delta,P_+)$ are shown schematically in Fig. 3.

We focus on the segments of the standard integer grid that are contained in $R(N,\Delta)$ and $A(N,\Delta,P_+)$, and their asymptotic behavior when $N \to \infty$. We note that the Euclidean area $\mathcal{A}$ of $R(N,\Delta)$ is

$$\mathcal{A}(R(N,\Delta)) = 2\pi\,\Delta\sqrt{N} + O(\Delta^2)$$

and the Euclidean area of $A(N,\Delta,P_+)$ is

$$\mathcal{A}(A(N,\Delta,P_+)) = 4/3\,\Delta\sqrt{2\Delta}N^{1/4} + O(\Delta^{5/2}N^{-1/4}).$$

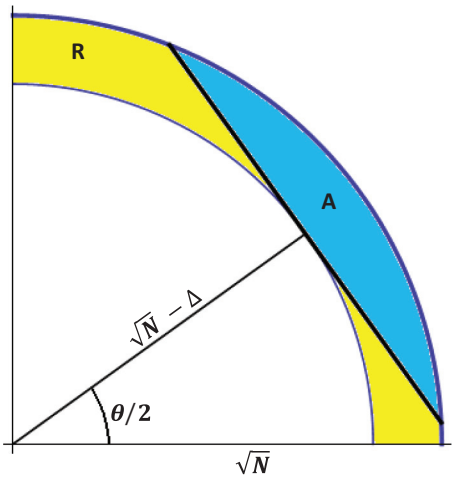

FIG. 3. (Color online) The ring $R(N,\Delta)$ (yellow, middle) and the segment $A(N,\Delta,P_+)$ (upper, blue, meniscus shaped), illustrated for values $N = 625$ and $\Delta = 4$.

Estimation of the number of integer grid points inside a flat contour is a known open problem with a rich history [18]. For our purposes, it suffices to know that the number of integer grid points

$$\{x,y \in \mathbb{Z}, (x,y) \in R(N,\Delta)\}$$

is asymptotically equal to $\Theta(\Delta\sqrt{N})$ and that the number of integer grid points

$$\{x,y \in \mathbb{Z}, (x,y) \in A(N,\Delta,P_+)\}$$

is asymptotically equal to $\Theta(\Delta^{3/2}N^{1/4})$. These claims can be proven by elementary geometric means.

Finally, we assume that $N = p^L$, where p is a fixed integer prime number with $p = 1 \mod 4$ and L is a large integer. Consider the set

$$s_4(N) = \{(x,y,z,w) \in \mathbb{Z}^4 \mid x^2 + y^2 + z^2 + w^2 = N\}$$

of all representations of N as a sum of squares of four integers. For $N = p^L$, the cardinality of the set is

$$\text{card}(s_4(N)) = 8(p^{L+1} - 1)/(p - 1) = \Omega(N).$$

This is an immediate consequence of the formula expressing $|s_4(N)|$ as eight times the sum of divisors of N (see [19]).

The projection $\text{Pr}_{x,y}(s_4(N))$ of $s_4(N)$ onto the first two coordinates consists of integer grid points (x,y) in the origin-centered circle of radius $\sqrt{N}$ satisfying a constraint $x^2 + y^2 = N - a^2 - b^2$ where (a,b) is some other grid point (a,b) (a witness) in the same circle. All witnesses (a,b) for (x,y) are on the same origin-centered circumference of radius $\sqrt{N - x^2 - y^2}$; the number of such witnesses is the number $r(N - x^2 - y^2)$ of the decompositions of $N - x^2 - y^2$ into sums of two squares. The average $\frac{1}{n}[r(1) + r(2) + \cdots + r(n)]$ converges to π [19,20]. Taking into account that the number of integers $\leqslant n$ decomposable into a sum of two squares is $\Theta(n/\sqrt{\log n})$ [21], we have that on average a point (x,y) has $O(\sqrt{L})$ witnesses, and so the cardinality of the projection is $\Omega(N/\sqrt{L})$ and the density in the circle is $\Omega(1/\sqrt{L})$. The essence of Conjecture 1 is that the density of the projection points in the ring and meniscus is no less than in the circle.

The conjecture is motivated by Ref. [22], although we count grid points (x,y) while Hooley counts corresponding integers x^2+y^2. The conjecture is presented for a general prime $p=1 \mod 4$, but our algorithms are developed for $p=5$, thus we need the conjecture to be true only for $p=5$.

Conjecture 1. Consider $N=p^L$, where p is a fixed integer prime number with $p=1 \mod 4$ and L is a large even integer. For a constant $\Delta>1$, let the four-square decomposition set $s_4(N)$, the geometric ring $R(N,\Delta)$, and the circular segment $A(N,\Delta,P_+)$ be defined as above. Let $\mathrm{Pr}_{x,y}(s_4(N))$ be the projection of the $s_4(N)$ onto its first two coordinates. Then

(1) $\mathrm{card}(\mathrm{Pr}_{x,y}(s_4(N)) \bigcap R(N,\Delta)) = \Omega(p^{L/2}/\sqrt{L})$,

(2) $\mathrm{card}(\mathrm{Pr}_{x,y}(s_4(N)) \bigcap A(N,\Delta,P_+)) = \Omega(p^{L/4}/\sqrt{L})$.

To empirically support the conjecture, first define the set

$$\mathrm{sn}(N,\Delta) = \{a^2+b^2 \mid a,b \in \mathbb{Z}, (a,b) \in R(N,\Delta)\}.$$

It is easy to see that

$$\mathrm{sn}(N,\Delta) \subset [p^L - 2\Delta\, p^{L/2}, p^L].$$

If $2\Delta\, p^{L/2} < p^L$, then the conditions of Theorem 1 in Ref. [22] are satisfied and the corollary implies that the cardinality of the set is

$$\mathrm{card}(\mathrm{sn}(N,\Delta)) = \Omega\left(\frac{p^{L/2}}{\sqrt{\log(p^L)}}\right) = \Omega\left(\frac{p^{L/2}}{L^{1/2}}\right).$$

Thus, there are as many distinct circumferences in $R(N,\Delta)$ that contain integer grid points, implying that the number of integer grid points on any one of these circumferences is $\Omega(L^{1/2})$ on average.

We further note that the set

$$v(L) = \{p^L - a^2 - b^2 \mid a,b \in \mathbb{Z}, (a,b) \in R(N,\Delta)\}$$

has cardinality

$$m = \mathrm{card}(v(L)) = \Omega\left(\frac{p^{L/2}}{L^{1/2}}\right).$$

Values from $v(L)$ are contained in the interval $[0, 2\,\Delta\, p^{L/2}]$. The average density of integers in that segment that are representable as a sum of two squares of integers is $\Omega[\sqrt{\log(N)}] = \Omega(\sqrt{L})$ [21]. Assuming that the density of such integers across the set $v(L)$ is the same, we infer from the assumption that there are at least $\Omega(p^{L/4}/\sqrt{L})$ integer grid points $(a,b) \in A(N,\Delta,P_+)$ that are projections of some four square decomposition of p^L (i.e., such that there exist $c,d \in \mathbb{Z}$ with $p^L = a^2+b^2+c^2+d^2$).

To verify statement (1) of Conjecture 1, we ran extensive computer simulations for $p=5$ and $L=\{16,\ldots,28\}$, and for $p=13$ and $L=\{12,\ldots,18\}$, using MATHEMATICA infinite precision integer arithmetic and observed behavior consistent with the conjecture. To motivate statement (2) of Conjecture 1, we tested the polar angles of points in $\mathrm{Pr}_{x,y}(s_4(N))$ for uniformity. The simulation covered $N = 5^{16},\ldots,5^{28},13^{12},\ldots,13^{18}$ and tested the null hypothesis that the distribution of the polar angles is uniform. Based on Kolmogorov-Smirnov statistics, the null hypothesis could not be rejected at any meaningful level of significance.

B. Algorithm

We now present the expected-polynomial time algorithm. We begin by approximating an arbitrary Z rotation with a $\langle \mathcal{C}+V\rangle$ circuit.

Problem 1. Given a Z-rotation $G=R_Z(\theta)$ and a small enough[1] target precision ϵ, synthesize a $\langle \mathcal{C}+V\rangle$ circuit $c(G,\epsilon)$ such that

$$\mathrm{dist}(c(G,\epsilon),G) < \epsilon \tag{4}$$

and

$$V_c(c(G,\epsilon)) \leqslant 4\,\log_5(2/\epsilon). \tag{5}$$

Theorem 3. There exists a randomized algorithm that solves Problem 1 in expected time polynomial in $\log(1/\epsilon)$.

We first present geometry that relates Theorem 3 to Conjecture 1 for $p=5$. Our goal is to select the target circuit depth value L such that

$$\epsilon < 2\times 5^{-L/4}. \tag{6}$$

Having found the smallest integer L satisfying Eq. (6), we then represent G as $G=\cos(\frac{\theta}{2})\,I + i\,\sin(\frac{\theta}{2})Z$ and consider approximating it with

$$U = (a\,I + b\,i\,X + c\,i\,Y + d\,i\,Z)\,5^{-L/2},$$

as suggested by Theorem 1.

Approximating G to precision ϵ in the trace distance metric is equivalent to finding U such that

$$\left[a\cos\left(\frac{\theta}{2}\right) + d\sin\left(\frac{\theta}{2}\right)\right]5^{-L/2} > 1-\epsilon^2.$$

For convenience we note that, without loss of generality, it suffices to prove the theorem for $-\pi/2 < \theta < \pi/2$ since we can always rotate the target gate to a position within this interval using $R_Z(\pm\pi/2)$ rotations from the Clifford group. We also note that our selection of L ensures that $5^{L/4}\,\epsilon \sim 2$.

Denote by $A_\epsilon(\theta)$ the segment of the unit disk where $[x\cos(\frac{\theta}{2}) + y\sin(\frac{\theta}{2})] > 1-\epsilon^2$. Let $D(L)$ be an isotropic dilation of the plane with coefficient $5^{L/2}$. Then the area of $D(L)[A_\epsilon(\theta)]$ is

$$\mathcal{A}(D(L)[A_\epsilon(\theta)]) = 5^L\frac{4\sqrt{2}}{3}\epsilon^3 \sim 8\frac{4\sqrt{2}}{3}5^{L/4}.$$

Define the angle $\phi = \sqrt{2}\epsilon(1-\epsilon^2/4)$ and the interval

$$I_w(\epsilon,\theta) = \left(5^{L/2}\sin\left(\frac{\theta}{2}-\phi\right), 5^{L/2}\sin\left(\frac{\theta}{2}+\phi\right)\right)$$

with subinterval $(5^{L/2}\sin(\frac{\theta}{2}-\epsilon), 5^{L/2}\sin(\frac{\theta}{2}+\epsilon))$. The length of the latter is approximately $2\times 5^{L/2}\cos(\frac{\theta}{2})\epsilon \geqslant 2\sqrt{2}\times 5^{L/4}$ and it contains approximately at least as many integer values.

Given any integer a such that

$$5^{L/2}\sin\left(\frac{\theta}{2}-\epsilon\right) < a < 5^{L/2}\sin\left(\frac{\theta}{2}+\epsilon\right),$$

[1] Although we do not have a closed-form bound on how small ϵ should be, our algorithm works well in practice for $\epsilon < 2\times 5^{-4} = 0.0032$.

we derive geometrically that the intersection of the horizontal line $w = a$ with $D(L)[A_\epsilon(\theta)]$ is a straight line segment that is longer than $5^{L/2}\frac{\epsilon^2}{2} \geqslant 2$ and that it contains at least two integer grid points. We are now ready to prove the theorem.

Proof. Revisiting the notations of the previous subsection, we introduce the set of all representations of 5^L as a sum of squares of four integers,

$$s_4(5^L) = \{(x,y,z,w) \in \mathbb{Z}^4 \mid x^2 + y^2 + z^2 + w^2 = 5^L\}.$$

The key step in the algorithmic proof is finding a point (a,d) in the intersection of $\Pr_{x,y}(s_4(5^L))$ and $D(L)[A_\epsilon(\theta)]$. Once such a point is found we can use a Rabin-Shallit algorithm [23] to express $5^L - a^2 - d^2$ as $b^2 + c^2, b, c \in \mathbb{Z}$. Then $U = (a\,I + b\,i\,X + c\,i\,Y + d\,i\,Z)\,5^{-L/2}$ would be the desired approximation of G, and can be represented precisely as a W circuit in at most L quaternion division steps.

Consider a horizontal line $w = a$, where $a \in I_w(\epsilon,\theta)$. By simple geometric calculation we find that the intersection of this line with the $D(L)[A_\epsilon(\theta)]$ segment is a line segment that is at most $5^{L/2}\epsilon^2/\cos(\frac{\theta}{2}) \leqslant 5^{L/2}\sqrt{2}\,\epsilon^2$ long.

For our choice of L this maximum length is approximately $4\sqrt{2}$ and thus the line segment contains at most five points with an integer first coordinate. On the other hand, we have shown earlier that for

$$5^{L/2}\sin\left(\frac{\theta}{2} - \epsilon\right) < a < 5^{L/2}\sin\left(\frac{\theta}{2} + \epsilon\right)$$

the intersection of the $w = a$ line with $D(L)[A_\epsilon(\theta)]$ is a line segment that is longer than 2 and must contain at least two points with integer z coordinate. In other words, if $a \in I_w(\epsilon,\theta)$ is a randomly selected integer, then with probability at least $1/\sqrt{2}$ the intersection segment contains at least two integer grid points.

Algorithm 2 gives the randomized approximation algorithm.

Algorithm 2 Randomized approximation

Require: Accuracy ϵ, angle θ
completion $\leftarrow$ null
$Sw \leftarrow$ set of all integers in $I_w(\epsilon,\theta)$
while completion == null **and** $Sw \neq \emptyset$ **do**
  Randomly, pick an integer a from Sw
  $Sw \leftarrow Sw - \{a\}$
  for all integer d such that $(d,a) \in D(L)[A_\epsilon(\theta)]$ **do**
    if exist $b,c \in \mathbb{Z}$ such that $5^L - a^2 - d^2 = b^2 + c^2$ **then**
      completion $\leftarrow (b,c)$
      Break;
    end if
  end for
end while
if completion==null **then**
  return null;
end if
$b \leftarrow$ first(completion)
$c \leftarrow$ last(completion)
return $U = (a\,I + b\,i\,X + c\,i\,Y + d\,i\,Z)\,5^{-L/2}$

In the worst case the algorithm terminates by exhausting the $\Theta(5^{L/4})$ candidate points in the $D(L)[A_\epsilon(\theta)]$ segment. However, we note that this segment is that of Conjecture 1 when $p = 5$. Therefore the share of satisfactory candidates among all of the integer grid points $D(L)[A_\epsilon(\theta)]$ is $\Omega(1/\sqrt{L})$. Thus the algorithm will terminate in $O(\sqrt{L})$ iterations on average.

Since the average overall number of iterations is moderate, the largest cost in the algorithm is line 7. It has been shown by Rabin and Shallit [23] that the effective test for an integer v to be a sum of squares of two integers has expected running cost of $O[\log^2(v)\,\log(\log(v))]$. In our case $v \leqslant 8 \times 5^{L/2}$ and we estimate the expected cost of the step as $O(L^2\log(L))$. Therefore the overall expected cost of the algorithm is $O(L^{5/2}\log(L))$ which translates into $O[\log(1/\epsilon)^{5/2}\log(\log(1/\epsilon))]$. ■

Remarkably, the above solution requires simpler number theory techniques than those developed for the $\{H,T\}$ basis in [4,5]. Conceptually, the simplification emerges for the same reasons that working with a transcendental extension of a numeric domain is typically easier than working with an algebraic extension. At their core, the methods in both [4] and [5] have algorithms for solving a norm equation in certain rings of algebraic numbers. While Selinger [4] constructs his own algorithm for cyclotomic numbers, Kliuchnikov *et al.* [5] use a method based on S units [24] as implemented in the PARI/GP software package. One can say that we also are solving a norm equation in line 7 of the above algorithm, however this is a norm equation over the simple domain of Gaussian integers and happens to nicely coincide with the well-studied classical problem of decomposing an integer into a sum of two squares [23].

C. Empirical tightening of the algorithm

As will be discovered in Sec. IV, most, although not all, single-qubit target gates allow approximating circuits with V count $V_c \leqslant 3\log_5(1/\epsilon)$ for small enough target precision ϵ. Thus the resource count $V_c = 4\log_5(1/\epsilon)$ guaranteed by Algorithm 2 is too loose in most cases, including in the case of most axial rotations. We thus consider tightening the algorithm by setting $L = \lceil c\log_5(1/\epsilon)\rceil$, $c < 4$ in an attempt to find a suitable completion candidate following the same randomized search loop as described in the previous subsection.

We iterate over c from 3 to 3.9 with a suitable step size δ_c. Since polynomial-time termination of the randomized search loop is no longer supported by Conjecture 1, we set an arbitrary time limit on the search loop for each iteration. The results of this algorithm on a test set of 1000 random rotations at six values of ϵ indicate that the minimal successful coefficient c satisfies $3.5 < c < 4$, with the average successful c estimated at 3.7. Figure 4 shows results for V count versus precision ϵ using the empirical tightening.

Example. The following decomposition approximates a rotation around the z axis by 0.1 radians, $R_Z(0.1)$, and has a V count of 76 at precision 4.3×10^{-15}: $V_1^{-1}\,V_2\,V_1\,V_2\,V_1^{-1}\,V_3\,V_2\,V_3\,V_1\,V_2\,V_3\,V_2^{-1}\,V_3^{-1}\,V_1V_2^{-1}\,V_3\,V_2\,V_3^{-1}\,V_1\,V_2^{-1}V_1^{-1}\,V_3\,V_1^{-1}\,V_3^{-1}\,V_2\,V_2\,V_2\,V_3^{-1}V_1^{-1}\,V_2\,V_1^{-1}\,V_2\,V_1^{-1}\,V_3^{-1}\,V_1\,V_2\,V_2\,V_3^{-1}V_2^{-1}\,V_2^{-1}\,V_1^{-1}\,V_2\,V_3^{-1}\,V_1^{-1}\,V_2\,V_2V_3\,V_1\,V_3\,V_1^{-1}\,V_3^{-1}\,V_1\,V_1$

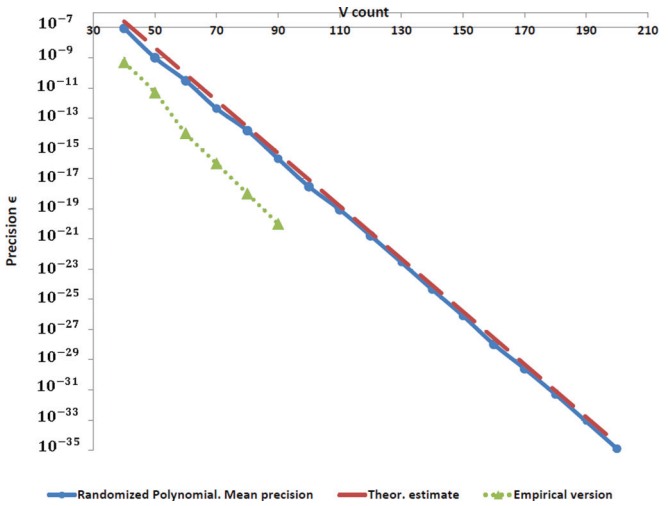

FIG. 4. (Color online) V count versus mean precision ϵ (measured in trace distance). Results are presented for 1000 random axial rotations for 17 values of the V count V_c. Dashed upper (red) line: theoretical bound on precision, $2 \times 5^{-V_c/4}$. Solid middle (blue) line: interpolated average precision. Dotted lower (green) line: average precision for an empirical version of the algorithm. Marker sizes are proportional to the standard deviations of the precision at each V count.

$V_3\ V_2^{-1}\ V_2^{-1}V_2^{-1}\ V_2^{-1}\ V_1^{-1}\ V_1^{-1}\ V_2\ V_3^{-1}V_1\ V_3^{-1}\ V_2^{-1}\ V_3\ V_1^{-1}$ $V_2\ V_1\ V_3^{-1}\ V_2\ V_2\ V_1^{-1}\ V_2\ V_2\ V_1$. This solution corresponds to $c = 3.7$.

D. Experimental results

We implemented Algorithm 2 from Theorem 3 in MATHEMATICA. Our implementation has the following simplifications:

(i) Line 7 has been redefined to return `PrimeQ`$[5^L - a^2 - d^2]$ for even a and d and to return false otherwise.[2]

(ii) Given a desired V count V_c, the algorithm terminates whenever a random candidate at distance less than $2 \times 5^{-V_c/4}$ from the target is picked.

We implicitly used the Rabin primality test since it is in general faster than complete integer factorization. We ran our MATHEMATICA solution over a set of 1000 random axial unitary rotations at 17 different V_c values. The test statistics are presented in Fig. 4. The solid blue line represents the interpolated average precision achieved over the test set. The sizes of the markers are proportional to the standard deviations of the precision at each level. The dashed red line shows the theoretical precision bound of $2 \times 5^{-V_c/4}$. The dotted green line charts the performance of the empirical version of the algorithm described in the previous subsection. Note that the tight match between the theoretical estimate and experimental results is not very insightful since the algorithm has been designed to terminate as soon as the theoretical precision has been achieved.

[2]`PrimeQ` is MATHEMATICA primality test that does not require complete factorization of the integer being tested. MATHEMATICA is a registered trademark of Wolfram Research, Inc.

The randomized algorithm can be used for approximate decomposition of any single-qubit unitary into a $\langle \mathcal{C} + V \rangle$ circuit since any $G \in \mathrm{SU}(2)$ can be decomposed exactly into three axial rotations, and the algorithm can be applied to each axial component. The V count in this case will scale as

$$V_c \leqslant 12 \log_5(6/\epsilon) = 12 \log_5(1/\epsilon) + 12 \log_5(6). \qquad (7)$$

However, for the majority of unitary gates, we can significantly reduce the depth of the output circuit with a corresponding increase in compilation time. The V count estimate given in Eq. (7) reflects the tripling of the circuit depth due to decomposition of the target unitary into three axial rotations. An alternative approach would be to perform a direct search in the four-dimensional integer grid; this will be the basis for our second algorithm described in Sec. IV.

E. Possible generalization

A foundation for efficient circuit synthesis over the V basis is the set of quaternions of norm 5^L and a body of number theory facts and conjectures related to that set. Given an integer prime p such that $p = 1 \mod 4$, it is apparent that most of these facts and observations generalize to quaternions of norm p^L, which are generated by the primitive ones of norm p. Modulo Lipschitz units there are $p + 1$ such quaternions in the generator set. These correspond to a basis of $p + 1$ unitary operators that we denote $V(p)$. Together with the Pauli gates a subset of $(p + 1)/2$ of the $V(p)$ operators generate the generalization of the W circuits.

However, in the case of $p = 5$, it was sufficient to add only one V operator in order to ensure the asymptotic uniformity of the grid of $\langle \mathcal{C} + \{V\} \rangle$ circuits. For $p > 5$, additional independent $V(p)$ operators are required.

For example, when $p = 13$ the following gates are required, in addition to the Clifford gates:

$$V_1(13) = (2\,I + 3\,i\,Z)/\sqrt{13},$$
$$V_2(13) = (I + 2\,i(X + Y + Z))/\sqrt{13},$$
$$V_3(13) = (2\,I + i\,X + 2\,i(Y + Z))/\sqrt{13},$$
$$V_4(13) = (2\,I + i\,Y + 2\,i(X + Z))/\sqrt{13},$$
$$V_5(13) = (2\,I + i\,Z + 2\,i(X + Y))/\sqrt{13}.$$

A generalization of Theorem 2 characterizes the gates representable exactly in the $\langle \mathcal{C} + \{V(p)\} \rangle$ basis as normalizations of Lipschitz quaternions of norm $p^L, L \in \mathbb{Z}$. The exact synthesis of the corresponding circuit for a unitary of the form

$$U = (a\,I + b\,i\,X + c\,i\,Y + d\,i\,Z)/p^{L/2}$$

amounts to a generalization of Algorithm 1 and requires at most $(p + 1) * L$ quaternion divisions. Theorem 3 also generalizes to $\langle \mathcal{C} + \{V(p)\} \rangle$ circuits, to the extent that Conjecture 1 holds for the prime parameter p, and the circuit depth estimate from the theorem generalizes to an estimate of the form $L \leqslant 4 \log_p(2/\epsilon)$.

We have chosen to focus on the $V(5)$ case for two reasons. First, the basis requires only one non-Clifford gate for which

we have a fault-tolerant implementation protocol. Second, we have so far only collected empirical data for $p = 5$.

IV. DIRECT SEARCH APPROXIMATION ALGORITHM

In this section, we present an algorithm based on optimized brute-force search for decomposing single-qubit unitaries into a circuit in the set $\langle \mathcal{P} + V \rangle$, where $\mathcal{P}$ is the set of single-qubit Pauli gates and V is one of the V gates. We first present relevant background.

A. Vicinity of a unitary in PSU(2) as a spherical cap

We begin by characterizing an ϵ neighborhood of a single qubit unitary as a "spherical cap" in a three-dimensional sphere S^3, i.e., as a portion of the sphere to one side of a certain three-dimensional hyperplane in the four-dimensional Euclidean space. Consider the four-dimensional Euclidean space with standard coordinates $\alpha,\beta,\gamma,\delta$. Let

$$S^3(R) = \{(\alpha,\beta,\gamma,\delta) \mid \alpha^2 + \beta^2 + \gamma^2 + \delta^2 = R^2\}$$

be the three-dimensional sphere of radius R centered at the origin. For any point on $S^3(R)$ we generate the unitary

$$\nu(\alpha,\beta,\gamma,\delta) = (\alpha\, I + i\,\beta\, X + i\,\gamma\, Y + i\,\delta\, Z)/R \in \mathrm{SU}(2).$$

The quantum gate group PSU(2) is the central quotient of SU(2) with the exact sequence $1 \to \mathbb{Z}_2 \to \mathrm{SU}(2) \to \mathrm{PSU}(2) \to 1$, therefore ν defines a $\mathbb{Z}_2$ covering of PSU(2) (which is the same factorization that is commonly used to glue an S^3 into three-dimensional projective space).

Under this covering the PSU(2) unitaries with nonzero trace are in one to one correspondence with the "northern" hemisphere,

$$S^3_+(R) = \{(\alpha,\beta,\gamma,\delta) \in S^3(R) | \alpha > 0\}.$$

Thus given a gate G, $|tr(G)| > 0$, then a small enough ϵ vicinity of that gate,

$$c_\epsilon(G) = \{U \in \mathrm{PSU}(2) | \ \mathrm{dist}(U,G) < \epsilon\},$$

is unambiguously identified with a spherical cap in $S^3_+(R)$.

To clarify, consider $G = \nu(\alpha,\beta,\gamma,\delta)$ and define

$$C_\epsilon(G) = \{(\alpha',\beta',\gamma',\delta') \in S^3_+(R) | \alpha\,\alpha' + \beta\,\beta' + \gamma\,\gamma' + \delta\,\delta' > R(1-\epsilon^2)\}.$$

Then $C_\epsilon(G)$ is a portion of $S^3(R)$ bounded by the hyperplane

$$\alpha\,\alpha' + \beta\,\beta' + \gamma\,\gamma' + \delta\,\delta' = R(1-\epsilon^2)$$

and $\nu(C_\epsilon(G)) = c_\epsilon(G)$.

We will focus further calculations on the ϵ neighborhoods that do not contain zero-trace gates and thus correspond to spherical caps completely contained in $S^3_+(R)$. It is trivial to modify all of the equations to cases where an ϵ neighborhood intersects the zero-trace "equator."

Given a $C_\epsilon(G)$ that is completely contained in $S^3_+(R)$, it is easy to derive, geometrically, that the metric volume $\mathcal{V}$ of $C_\epsilon(G)$ is

$$\mathcal{V}(C_\epsilon(G)) = 4\pi R^3 \int_0^{\cos^{-1}(\epsilon')} \sin^2(\eta) d\eta = 2\pi R^3 \left(\cos^{-1}(\epsilon') - \frac{1}{2}\sin[2\cos^{-1}(\epsilon')]\right),$$

where $\epsilon' = 1 - \epsilon^2$. Taking the Taylor series expansion of the latter at $\epsilon = 0$, we find that

$$\mathcal{V}(C_\epsilon(G)) = \frac{8\pi\sqrt{2}\epsilon^3 R^3}{3} + O(\epsilon^5).$$

In the next sections we focus on precision targets ϵ for which the $C_\epsilon(G)$ neighborhoods have sufficient metric volume.

B. Bound for uniform precision

We start by establishing that there exist unitary gates in PSU(2) that cannot be approximated by W circuits of $V_c \leqslant L$ to a precision better than $\epsilon_L = 5^{-L/4}/2$. This is based on the following observation:

Observation 1. Let w be a W circuit different from the identity with $V_c(w) \leqslant L$, then it evaluates to $U(w)$ with $|\mathrm{tr}(U(w))| \leqslant 2(1 - 5^{-L/2})$ and $U(w)$ is at least $5^{-L/4}$ away from the identity.

Indeed

$$U(w) = \frac{a}{5^{L/2}} I + \frac{i\,(b\,X + c\,Y + d\,Z)}{5^{L/2}},\ a,b,c,d \in \mathbb{Z}.$$

Since $U(w)$ is not the identity, $|a|$ cannot be greater than $5^{L/2} - 1$.

Now, let $P \in \{I,X,Y,Z\}$ be a Pauli gate.

Observation 2. A circuit w with $V_c(w) \leqslant L$ and distinct from P evaluates to $U(w)$ with a distance at least $5^{-L/4}$ from P.

Indeed, if w is a W circuit at a certain distance from P then wP is a circuit with the same V count at the same distance from the identity. Thus, if $\epsilon < \epsilon_L = 5^{-L/4}/2$ and $G \in \mathrm{PSU}(2)$ is any unitary such that

$$\epsilon < \mathrm{dist}(G,P) < 2\,\epsilon_L - \epsilon,$$

then there are no W circuits of $V_c \leqslant L$ within distance ϵ from G by the triangle inequality for dist:

$$\forall w,\ \mathrm{dist}(w,G) \geqslant \mathrm{dist}(w,P) - \mathrm{dist}(G,P) > \epsilon.$$

On the other hand, $\mathrm{dist}(G,P)$ is also greater than ϵ. Therefore, the uniform precision guarantee cannot be better than $5^{-L/4}/2$ for W circuits of $V_c \leqslant L$. In other words, the uniform guarantee of optimal circuit depth cannot be better than $4\log_5(1/\epsilon) - 4\log_5(2)$.

Revisiting the above discussions, we note that for $\epsilon < \epsilon_L = 5^{-L/4}/2$ there exist "exclusion zones" of width $2(\epsilon_L - \epsilon)$ around each of the Pauli gates consisting of unitaries that cannot be approximated to precision ϵ by W circuits with $V_c \leqslant L$. Using the spherical cap volume formulas from the previous subsection, for ϵ significantly smaller than ϵ_L, we estimate the combined volume of these exclusion zones, relative to the volume of the S^3_+ as $O[5^{-L/2}(5^{-L/4} - 3\,\epsilon)]$.

C. A working conjecture

Given the set of W circuits with $V_c \leqslant L$, we will consider two key precision targets: $\epsilon_4(L) = 2 \times 5^{-L/4}$ and $\epsilon_3(L) = 5^{-L/3}$. Consider the three-dimensional hemisphere $S^3_+(5^{L/2})$. As per the results from the previous subsection, the metric volumes of the ϵ_4 and ϵ_3 neighborhoods are

$$\mathcal{V}(C_{\epsilon_4(L)}(G)) \sim \frac{64\pi\sqrt{2}\,5^{3L/4}}{3}$$

and

$$\mathcal{V}(C_{\epsilon_3(L)}(G)) \sim \frac{8\pi\sqrt{2}\,5^{L/2}}{3}.$$

Since the volume of $S^3_+(5^{L/2})$ is equal to $\pi^2 5^{3L/2}$, the relative metric share that these neighborhoods occupy on the hemisphere are

$$\frac{\mathcal{V}(C_{\epsilon_4(L)}(G))}{\mathcal{V}(S^3_+(5^{L/2}))} \sim \frac{64\sqrt{2}\,5^{-3L/4}}{3\pi}$$

and

$$\frac{\mathcal{V}(C_{\epsilon_3(L)}(G))}{\mathcal{V}(S^3_+(5^{L/2}))} \sim \frac{8\sqrt{2}\,5^{-L}}{3\pi},$$

respectively.

Conjecture 2. (1) There exists a positive integer L_4 such that for any integer $L > L_4$ and any single-qubit gate G there exists a Wcircuit w such that

$$\mathrm{dist}(G,w) \leqslant \epsilon_4(L).$$

(2) For large enough integer L ($L > L_3$) there exists an open subset $\mathbb{G}_3 \subset \mathrm{PSU}(2)$ with metric volume $[1 - o(1)]\mathcal{V}(S^3_+)$ (when $L \to \infty$) such that for each $G \in \mathbb{G}_3$ there exists a W circuit w with

$$\mathrm{dist}(G,w) \leqslant \epsilon_3(L).$$

The common motivation for both clauses of this conjecture is that the number of distinct W circuits scales as 5^{V_c}. More specifically, there are approximately 5×5^L distinct unitaries in PSU(2) that are represented exactly by W circuits with $V_c \leqslant L$.

This stems from the fact that 5^L has exactly $10(5^L - 2)$ distinct decompositions into a sum of four squares of integers, which can be easily derived from the Jacobi formula for the r_4 function:

$$r_4(n) = 8 \sum_{SC(d)} d, SC(d) = (d \mid n)\ \&\ (d \mod 4 \neq 0)$$

[see chapters on the $r(n)$ function in Ref. [19]]. Geometrically, there are exactly $10(5^L - 2)$ distinct integer grid points on $S^3(5^{L/2})$ and the set of such grid points is central-symmetrical with respect to the origin, so approximately half of these integer grid points lie on the $S^3_+(5^{L/2})$ piece of the hemisphere. Further intuition in support of the conjectures is drawn from [1,2], which investigate the distribution density of the elements of the free group generated by $\langle V_1, V_2, V_3, V_1^{-1}, V_2^{-1}, V_3^{-1}\rangle$.

A stronger special case of Conjecture 2 postulates that for any

$$G = v(\alpha,\beta,\gamma,\delta) = (\alpha\, I + i\,\beta\, X + i\,\gamma\, Y + i\,\delta\, Z),$$

where $\alpha^2 + \beta^2 + \gamma^2 + \delta^2 = 1$, there is an integer grid point on $S^3(5^{L/2})$ within distance $\leqslant 2$ of $(\alpha,\beta,\gamma,\delta) \times 5^{L/2}$. Although we do not claim that this stronger statement is true for all unitaries G, the perceived near-uniformness of the distribution of the integer lattice grid points over $S^3(5^{L/2})$ for large enough L makes it plausible for most unitaries.

D. Algorithm outline

Our algorithm to address Problem 2 below employs optimized direct search.

Problem 2. Given an arbitrary single-qubit unitary $G \in \mathrm{PSU}(2)$ and a small enough target precision ϵ, synthesize a W circuit $c(G,\epsilon)$ such that

$$\mathrm{dist}(c(G,\epsilon),G) < \epsilon \tag{8}$$

and the V count of the resulting circuit is

$$V_c \leqslant 3\,\log_5(1/\epsilon) \tag{9}$$

for the majority of target unitaries and

$$V_c \leqslant 4\,\log_5(2/\epsilon) \tag{10}$$

in edge cases.

Let L be the intended V count of the desired approximation circuit. Given a target single-qubit unitary gate represented as $G = \alpha I + \beta i X + \gamma i Y + \delta i Z$, in order to find integers (a,b,c,d) such that $a^2 + b^2 + c^2 + d^2 = 5^L$ and

$$\mathrm{dist}(G,(a\,I + b\,i\,X + c\,i\,Y + d\,i\,Z)5^{-L/2}) < \epsilon, \tag{11}$$

we split the $\alpha,\beta,\gamma,\delta$ coordinates into two-variable blocks. Let us assume that the split is given by $(\alpha,\delta),(\beta,\gamma)$. For the approximation inequality in Eq. (11) to hold it is sufficient that

$$(b\,5^{-L/2} - \beta)^2 + (c\,5^{-L/2} - \gamma)^2 < \epsilon^2 \tag{12}$$

and

$$(a\,5^{-L/2} - \alpha)^2 + (d\,5^{-L/2} - \delta)^2 < \epsilon^2. \tag{13}$$

Our goal is to achieve $\epsilon = 5^{-L/3}$. It is easy to see that there are approximately $\pi 5^{L/3}$ integer pairs satisfying each of the conditions in Eqs. (12) and (13) for that ϵ. We can now sweep over all of the (b,c) integer pairs and build a hash table of all of the $5^L - b^2 - c^2$ differences occurring in the first set. Then we can sweep over all of the (a,d) integer pairs from the second set, in search of one for which $a^2 + d^2$ occurs in the hash table.

Using number-theoretical considerations (see, for example, [23]), one can reduce the number of candidates considered in this direct search by a factor of approximately $\frac{LR}{2\sqrt{L\,\ln(5)}}$ (where LR is the Landau-Ramanujan constant). Thus, for $L = 34$ the reduction factor is approximately 0.05.

For target unitaries that cannot be approximated to precision $5^{-L/3}$, the algorithm iteratively triples the precision goal (which has an effect of expanding the search space at each iteration) until the satisfactory candidate is found. The outline of the algorithm is given in Algorithm 3.

Algorithm 3 Direct search approximation

Require: Accuracy ϵ, Target gate $G = \alpha I + \beta i X + \gamma i Y + \delta i Z$

$L \leftarrow \lceil 3 \times \log_5(1/\epsilon) \rceil$
hash $\leftarrow$ Dictionary$\langle$Integer,(Integer $*$ Integer)$\rangle$
bound$\pm \leftarrow 5^L (\sqrt{\alpha^2 + \delta^2} \pm \epsilon)^2$
for all $b,c \in \mathbb{Z}$ satisfying Eq (12) **do**
  if bound$- \leqslant 5^L - b^2 - c^2 \leqslant$bound$+$ **and** $5^L - b^2 - c^2$ is decomposable into two squares **then**
    Add $(5^L - b^2 - c^2,(b,c)) \rightarrow$hash
  end if
end for
completion $\leftarrow$ fail
for all integer pairs (a,d) satisfying Eq (13) **do**
  if hash contains key equal to $a^2 + d^2$ **then**
    completion $\leftarrow (a,b,c,d)$
    Break;
  end if
end for
if completion $\neq$ fail **then**
  completion $\leftarrow$ completion.$(I,i\,X,i\,Y,i\,Z)5^{-L/2}$
end if
return completion

E. Experimental results and comparison

The chart in Fig. 5 presents the results of evaluating our direct search algorithm on a set of 1000 random unitaries. The vertical axis plots precision ϵ on a logarithmic scale. The horizontal axis plots the maximum V count allowed in the approximating circuit. The dashed pink curve represents the tight precision target of $5^{-V_c/3}$. The solid blue curve represents the average approximation distance over the set of test unitaries; the error bars measure the standard deviation around the average. The green dotted curve plots the worst cases. For a small number of test unitaries, the algorithm could not find an approximating sequence with precision $5^{-V_c/3}$ or better for V count V_c.

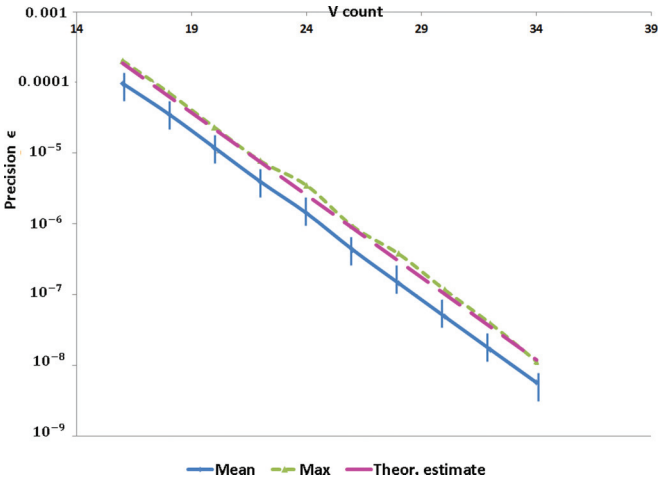

FIG. 5. (Color online) V count vs mean precision ϵ (measured by trace distance) of the approximation of 1000 random unitaries. The plot shows the $5^{-V_c/3}$ precision goal (dashed middle line, pink), the experimental average (solid lower line, blue), and the worst cases (dotted upper line, green).

TABLE I. V counts for precision ϵ for two V basis decomposition algorithms: randomized approximation (RA) (Algorithm 2) and direct search (DS) (Algorithm 3), for 1000 random nonaxial rotations. Columns 2 and 3 list the median V count; column 4 lists the V count for the worst case; column 5 gives the factor of improvement between RA median and DA median.

ϵ	RA, median	DS, median	DS, worst	Imp. factor
10^{-3}	56.5	13	15	4.35
10^{-4}	73.5	15.9	18	4.62
10^{-5}	91	20.5	22	4.44
10^{-6}	108	24.6	26	4.39
10^{-7}	125	28.95	31	4.32
10^{-8}	142.5	33.2	35	4.29
10^{-9}	159.5	37.3	39	4.28

In practice, experimental evidence suggests that this algorithm works well for the majority of nonaxial unitary rotations. We have found that approximation circuits obtained by Algorithm 2 are about four times deeper than the circuits produced by direct search using Algorithm 3. This factor primarily arises because the nonaxial rotation is first broken into three axial components and then a more liberal precision of $\epsilon_4(L)$ is pursued for each component.

Table I compares the V count and precision values for Algorithms 2 and 3, for precisions between 10^{-3} and 10^{-9}. Results indicate approximately a factor of 4 improvement in V count when using Algorithm 3. The improvement is also apparent from the plot shown in Fig. 1 (green versus black curves).

On a single desktop computer, a precision of 10^{-9} is approximately the limit for the direct search algorithm. At this precision, our implementation using .NET dictionaries requires up to 100 GB of memory and up to 1 h CPU time for each approximation target. Due to the exponential nature of the direct search algorithm, the same compilation requires only 3 sec and very small memory for a precision around 10^{-6}. In practice, for many algorithms, a precision of 10^{-6} is sufficient for approximating single-qubit unitaries [25]. It should also be noted that the search space of the algorithm can easily be tiled and distributed, which would allow precisions of 10^{-12} to be achieved when using a high performance cluster or cloud.

V. CONCLUSIONS AND FUTURE WORK

In conclusion, we have proposed two algorithms for decomposing a single-qubit unitary into the V basis, an efficiently universal basis that may have advantages over decomposing into the $\{H,T\}$ basis. Our algorithms produce efficient circuits that approximate a single-qubit unitary with high precision, and are computationally efficient in practice. Another application of our algorithm is to Toffoli-based circuits. To the best of our knowledge, an efficient constructive algorithm for decomposing into the basis containing Clifford gates plus the Toffoli gate has not been previously proposed. Our methods provide an algorithmic solution: first compile into

the V basis and subsequently replace each V by a Toffoli-based circuit (Fig. 2).

A key direction for future research is to determine a low-cost, exact implementation of a V gate, which could include a native implementation on a given quantum computer architecture. In Appendix, we present two possible exact constructions, however the cost of these constructions does not immediately merit decomposing into the V basis instead of the $\{H,T\}$ basis. Discovery of improved implementations could make decomposing into the V basis advantageous over decomposing into the $\{H,T\}$ basis. In order for decomposition into the V basis to be cost competitive with state-of-the-art $\{H,T\}$ decomposition, it is necessary to determine an exact, fault-tolerant V gate implementation that costs less than the cost of 6 T gates (where the complete cost of a T gate could be determined, for example, based on a given state distillation protocol [13–15]).

ACKNOWLEDGMENTS

We thank A. Blass for numerous discussions, C. Jones for the reference to the Toffoli-based implementation of the V gate, and G. Martin for his constructive critique of our number-theoretic conjectures.

APPENDIX: V GATE IMPLEMENTATION

A V gate can be *approximated* using, for example, a $\{H,T\}$ decomposition algorithm, but results in a circuit of length 70 or more, depending on the desired precision. Here, we describe two possible *exact* implementations of the V gate, in that they achieve perfect precision. We outline techniques for implementing the V_3 gate exactly:

$$V_3 = (I + 2\,i\,Z)/\sqrt{5}.$$

In matrix form, this gate can be represented as

$$V_3 = \frac{1}{\sqrt{5}}\begin{bmatrix} 1+2i & 0 \\ 0 & 1-2i \end{bmatrix} = \frac{1+2i}{\sqrt{5}}\begin{bmatrix} 1 & 0 \\ 0 & \frac{-(3+4i)}{5} \end{bmatrix}.$$

The other V gates can then be implemented using V_3 and Clifford gates.

1. Implementing V_3 with Toffolis

The first method uses the circuit given in Fig. 2, and is a slight modification of Fig. 4.17 in Ref. [9]. It requires two Toffoli gates and two ancillae qubits. With probability 5/8, V_3 is applied to the target qubit, otherwise the identity is applied. The expected number of times this circuit must be repeated in order to result in the application of V_3 is 8/5, with an expected cost of 3.2 Toffolis. With advances in fault-tolerant Toffoli gate implementations [10–12], this implementation of V_3 becomes appealing.

In addition, a generalization to $V_3(p_i)$ gates, where $p_0 = 5, p_1 = 17, p_2 = 101, p_3 = 257, p_4 = 4097, \ldots,$ can be implemented using a nested circuit based on Fig. 2, where the S gate is replaced with $V(p_{i-1})$ for $i > 0$. When $i = 0$, the circuit is the same as in Fig. 2. The expected number of applications of the outermost circuit rapidly approaches 1 as p_i grows. Table II gives the expectations for $p_0, \ldots, p_4$.

TABLE II. The expected number of applications of the outermost circuit for select values of p_i.

Gate	Expectation
$V_3(5)$	8/5
$V_3(17)$	20/17
$V_3(101)$	104/101
$V_3(257)$	260/257
$V_3(4097)$	4100/4097

2. Implementing V_3 with state distillation

The second method uses the protocol given in Ref. [8]. For additional details on the protocol, we refer the reader to Ref. [8].

Ignoring global phase, we first solve for the angle of rotation θ about the z axis invoked by V_3 using the following identity:

$$e^{i\theta} = \cos\theta + i\sin\theta = \frac{-(3+4i)}{5}$$
$$\Rightarrow \theta = \cos^{-1}\left(-\frac{3}{5}\right) \approx 4.068\,89.$$

Consider the angle $\theta' = \cos^{-1}(\frac{3}{5}) \approx 0.927\,295$. This angle is π away from θ: $\theta = \theta' + \pi$. Thus, if we want to implement rotation $R_Z(\theta)$, we can implement the gate sequence $R_Z(\theta) = R_Z(\theta')R_Z(\pi)$, where $R_Z(\pi)$ is the Pauli Z gate and

$$\theta' = 2\theta_2 + \frac{\pi}{4},$$

where $2\theta_2$ is the angle resulting from using the resource state $e^{-i\pi/8}HS^\dagger|H_2\rangle$. The $\frac{\pi}{4}$ part of the angle is a $T = R_Z(\pi/4)$ gate, thus $R_Z(\theta') = R_Z(2\theta_2)T$, and $R_Z(\theta) = R_Z(2\theta_2)TZ$.

The circuit to obtain a rotation of $Z(2\theta_2)$ is given in Fig. 6. The circuit results in the application of $\pm 2\theta_2$ to $|\psi\rangle$, each with equal probability. If $m = 0$, then $R_Z(2\theta_2)$ has been applied. If $m = 1$, we must apply $R_Z(4\theta_2)$. Further details on the $m = 1$ case are given in Sec. A 4.

3. Obtaining an $|H_2\rangle$ resource state

To implement V_3, we require a nonstabilizer state $|H_2\rangle$, which can be obtained using the ladder given in Ref. [8]. We begin by describing how to obtain the ladder state $|H_2\rangle$, and then describe how to implement V_3 using this resource state.

The circuit of Fig. 7 measures the parity of the two input qubits and decodes the resulting state into the second qubit. Let the two inputs be magic states $|H\rangle$ and define $\theta_0 = \frac{\pi}{8}$:

$$|H\rangle = |H_0\rangle = \cos\theta_0|0\rangle + \sin\theta_0|1\rangle.$$

Upon application of the controlled-NOT gate $\Lambda(X)$,

$$|H_0\rangle|H_0\rangle \xrightarrow{\Lambda(X)} \cos^2\theta_0|00\rangle + \sin^2\theta_0|01\rangle + \cos\theta_0\sin\theta_0(|11\rangle + |10\rangle).$$

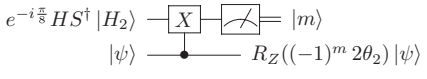

FIG. 6. Circuit to rotate by angle $\pm 2\theta_2$ around the z axis.

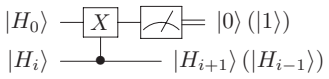

FIG. 7. Two-qubit circuit used to obtain $|H_i\rangle$ states from initial resource states $|H_0\rangle$. Upon measuring the 0 (1) outcome, the output state is $|H_{i+1}\rangle$ ($|H_{i-1}\rangle$).

Upon measurement m of the first qubit, we have

$$\begin{aligned} &\xrightarrow{m=0} \quad \frac{\cos^2\theta_0|0\rangle + \sin^2\theta_0|1\rangle}{\cos^4\theta_0 + \sin^4\theta_0}, \text{ or} \\ &\xrightarrow{m=1} \quad \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle). \end{aligned}$$

We define θ_1 such that

$$\cos\theta_1|0\rangle + \sin\theta_1|1\rangle = \frac{\cos\theta_0|0\rangle + \sin\theta_0|1\rangle}{\cos^4\theta_0 + \sin^4\theta_0},$$

from which we deduce $\cot\theta_1 = \cot^2\theta_0$.

Thus we have $|H_1\rangle = \cos\theta_1|0\rangle + \sin\theta_1|1\rangle$, a nonstabilizer state obtained from $|H\rangle$ states, Clifford operations, and measurements. If the measurement outcome is 1, then we obtain a stabilizer state and discard the output (see Fig. 7). The measurement outcomes occur with respective probabilities $p_{m=0,0} = \cos^4\theta_0 + \sin^4\theta_0 = \frac{3}{4}$ and $p_{m=1,0} = 1 - p_0 = \frac{1}{4}$.

Now consider the next step of the ladder. We recurse on this protocol using the nonstabilizer states produced by the previous round of the protocol as input to the circuit in Fig. 7. In this case, we need only go to state $|H_2\rangle$, which is defined as

$$|H_2\rangle = \cos\theta_2|0\rangle + \sin\theta_2|1\rangle,$$

where $\cot\theta_2 = \cot^3\theta_0$.

To obtain this state, we use as input the previously produced $|H_1\rangle$ state and a new $|H_0\rangle$ state:

$$\begin{aligned} |H_0\rangle|H_1\rangle \xrightarrow{\Lambda(X)} &\cos\theta_0\cos\theta_1|00\rangle + \sin\theta_0\sin\theta_1|01\rangle \\ &+ \sin\theta_0\cos\theta_1|10\rangle + \cos\theta_0\sin\theta_1|11\rangle. \end{aligned}$$

Upon measurement of the first qubit, we have

$$\begin{aligned} &\xrightarrow{m=0} (\cos\theta'|0\rangle + \sin\theta'|1\rangle), \\ &\xrightarrow{m=1} (\cos\theta''|0\rangle + \sin\theta''|1\rangle), \text{ where} \\ \cot\theta' &= \cot\theta_1\cot\theta_0 = \cot^3\theta_0 = \cot\theta_2, \\ \cot\theta'' &= \cot\theta_1\tan\theta_0 = \cot^1\theta_0 = \cot\theta_0. \end{aligned}$$

Thus, if we measure $m = 0$, we obtain the state $|H_2\rangle$ and if we measure $m = 1$, we obtain $|H_0\rangle$. The probability of measuring 0 is given by

$$p_{m=0,1} = \cos^2\theta_1\cos^2\theta_0 + \sin^2\theta_1\sin^2\theta_0.$$

Note that $\frac{3}{4} \leqslant p_{m=0,i} < \cos^2\frac{\pi}{8} = 0.853\ldots$, so the probability of obtaining $|H_2\rangle$ is far higher than the probability of obtaining $|H_0\rangle$.

4. Resource cost

What is the cost of obtaining a $|H_2\rangle$ state in terms of $|H_0\rangle$ resource states? We simulated 10 million instances of the ladder to determine the average cost of obtaining $|H_1\rangle$ and $|H_2\rangle$. Recall that the probabilities of moving "up" the ladder are higher than moving "down" the ladder. For $|H_1\rangle$, the cost is on average 2.66 $|H_0\rangle$ states, with a median cost of 2. For $|H_2\rangle$, the cost is on average 4.35 $|H_0\rangle$ states, with a median cost of 3.

What is the cost of implementing $R_Z(\theta')$? Recall that our technique uses a probabilistic circuit with a success probability of 1/2. Thus, on average it will require two attempts for success. If the circuit succeeds, the cost in $|H_0\rangle$ states is roughly 5.35. If the circuit fails, then we must correct the circuit by applying a Z rotation of $2 \times 2\theta_2$.

This requires preparing a resource state $R_Z(4\theta_2)$, which can be done using the circuit given in Fig. 6 with $|\psi\rangle = e^{-i\pi/8}HS^\dagger|H_2\rangle$. On average, two attempts will be required to prepare the state, resulting in an average cost of 4 $|H_2\rangle$ states, or roughly $4 \times 4.35 = 17.4$. The prepared state is applied to the target qubit $|\psi\rangle$ using the same circuit in Fig. 6, except now the top input qubit is $|R_Z(4\theta_2)\rangle$. The total cost if the circuit succeeds on this second attempt, after the first failure, is $1 + 4.35 + 17.4 = 22.75$.

As can be seen, each attempt that fails requires preparation of a more costly resource state for the next attempt. The series of attempts is a negative binomial of parameter $p = \frac{1}{2}$ and the expected number of attempts to achieve success goes as $\sim \frac{1}{p} = 2$. In general, at attempt k, a resource state to perform rotation by angle $2^k \times 2\theta_2$ is required. The cost of preparing the resource state grows exponentially in k, and in the limit is infinite. However, in practice, we will only make one to three attempts, and upon the final failure, apply a different approximation technique to the remaining rotation,[3] using methods of, for example, Refs. [4] and [5]. The optimal number of attempts to make before backing off to a different technique can be determined based on the required precision level (since the backoff method will only be approximate) and the chosen technique.

[3]We may in fact apply the backoff technique to the entire remaining sequence, that is, by determining the unitary from the remaining sequence and approximating it with the backoff technique.

[1] A. Lubotsky, R. Phillips, and P. Sarnak, Commun. Pure Appl. Math. **34**, 149 (1986).
[2] A. Lubotsky, R. Phillips, and P. Sarnak, Commun. Pure Appl. Math. **40**, 401 (1987).
[3] A. W. Harrow, B. Recht, and I. L. Chuang, J. Math. Phys. **43** (2002).
[4] P. Selinger, arXiv:1212.6253.
[5] V. Kliuchnikov, D. Maslov, and M. Mosca, arXiv:1212.6964.
[6] V. Kliuchnikov, D. Maslov, and M. Mosca, Quantum Inf. Comput. **13**, 607 (2013).
[7] V. Kliuchnikov, D. Maslov, and M. Mosca, Phys. Rev. Lett. **110**, 190502 (2013).
[8] G. Duclos-Cianci and K. M. Svore, arXiv:1210.1980.
[9] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, UK, 2000).

[10] B. Eastin, Phys. Rev. A **87**, 032321 (2013).
[11] C. Jones, Phys. Rev. A **87**, 022328 (2013).
[12] C. Jones, arXiv:1303.6971.
[13] S. Bravyi and A. Kitaev, Phys. Rev. A **71**, 022316 (2005).
[14] B. W. Reichardt, Quantum Inf. Process. **4**, 251 (2005).
[15] A. M. Meier, B. Eastin, and E. Knill, arXiv:1204.4221.
[16] R. Lipschitz, *Untersuchungen uber die Summen von Quadraten* (Max Cohen und Sohn, Bonn, 1886).
[17] J. Conway and D. Smith, *On Quaternions And Octonions: Their Geometry, Arithmetic, and Symmetry* (Peters, Natick, MA, 2003).
[18] R. Guy, *Unsolved Problems in Number Theory*, 3rd ed. (Springer, New York, 2004).
[19] G. Hardy and E. Wright, *An Introduction to the Theory of Numbers* (Claredon, Oxford, 1979).
[20] E. Kraetzel, *Lattice Points* (Kluwer, Dordrecht, 1988).
[21] E. Landau, Arch. der Math. und Phys. **13**, 305 (1908).
[22] C. Hooley, J. Reine Angew. Math. **267**, 207 (1974).
[23] M. Rabin and J. Shallit, Commun. Pure Appl. Math. **39**, 239 (1986).
[24] D. Simon, Math. Comput. **74**, 1531 (2005).
[25] A. G. Fowler and L. C. Hollenberg, Phys. Rev. A **70**, 032329 (2004).