

Efficient synthesis of probabilistic quantum circuits with fallback

Alex Bocharov, Martin Roetteler, and Krysta M. Svore

Quantum Architectures and Computation Group, Microsoft Research, Redmond, Washington, USA

(Received 16 March 2015; published 18 May 2015)

Repeat-until-success (RUS) circuits can approximate a given single-qubit unitary with an expected number of T gates of about $\frac{1}{3}$ of what is required by optimal, deterministic, ancilla-free decompositions over the Clifford + T gate set. In this work, we introduce a more general and conceptually simpler circuit decomposition method that allows for synthesis into protocols that probabilistically implement quantum circuits over several universal gate sets including, but not restricted to, the Clifford + T gate set. The protocol, which we call probabilistic quantum circuits with fallback (PQF), implements a walk on a discrete Markov chain in which the target unitary is an absorbing state and in which transitions are induced by multiqubit unitaries followed by measurements. In contrast to RUS protocols, the presented PQF protocols are guaranteed to terminate after a finite number of steps. Specifically, we apply our method to the Clifford + T , Clifford + V , and Clifford + $\pi/12$ gate sets to achieve decompositions with expected gate counts of $\log_b(1/\varepsilon) + O[\ln[\ln(1/\varepsilon)]]$, where b is a quantity related to the expansion property of the underlying universal gate set.

DOI: [10.1103/PhysRevA.91.052317](https://doi.org/10.1103/PhysRevA.91.052317)

PACS number(s): 03.67.Lx, 03.67.Ac

I. INTRODUCTION

Techniques to efficiently compile higher-level quantum algorithms into lower-level fault-tolerant circuits are a critical step for the implementation of a scalable, general purpose quantum computer. Several universal fault-tolerant gate sets arise from augmenting the set of Clifford gates by additional gates that arise naturally from the underlying fault-tolerance scheme. An important example is the Clifford + T basis, consisting of controlled-NOTs (CNOT) and Hadamard (H) gates, together with the T gate, which is given by $T = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{bmatrix}$. Further examples of interest are the Clifford + V basis in which the set of Clifford gates is augmented by the six non-Clifford gates $\frac{1}{\sqrt{5}}(1 \pm 2iP)$, where $P \in \{X, Y, Z\}$, and the Clifford + $\pi/12$ basis in which the gate $K = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/6} \end{bmatrix}$ is added.

While the Solovay-Kitaev algorithm [1,2] allows us to solve the synthesis problem for any universal gate set, there are certain disadvantages to this approach, in particular the large depth of the resulting circuits: to the best of our knowledge, the resulting depth is only known to scale as $O[\log_2^{3.97}(1/\varepsilon)]$, where ε is the target approximation error with which the single-qubit unitary has to be implemented. Also, the compilation time of the Solovay-Kitaev method, i.e., the time it takes to execute the classical algorithm that produces the output circuit is quite high, namely, almost cubic in $\ln(1/\varepsilon)$. This makes the application of the algorithm for small values of the target precision, say in a regime where $\varepsilon \sim 10^{-15}$, difficult if not impossible. On the other hand, there exist several quantum algorithms that would require this level of target approximation error in order to scale to instance sizes of practical interest.

Happily, it was shown recently [3–6] that for the Clifford + T basis, elementary number theory can be leveraged to obtain much more efficient algorithms for approximating a single-qubit gate. We refer to these methods as being deterministic and ancilla free as they lead to a decomposition of the target unitary that can be executed in an entirely predetermined sequence of single-qubit unitaries over the

given gate set. The number of T gates in the resulting circuits scales close to $3 \log_2(1/\varepsilon)$ for Z rotations, which is within a constant factor of the information-theoretic lower bound. Also, the compilation time of these methods is low: using reasonable number-theoretic conjectures for which there exists an overwhelming amount of empirical evidence, the compilation time follows the same scaling (up to logarithmic factors). Any nonaxial rotation V can be decomposed into axial rotations such that [1]

$$V = e^{i\delta} R_z(\alpha) H R_z(\beta) H R_z(\gamma) \quad (1)$$

for real values $\alpha, \beta, \gamma, \delta$. This yields an upper bound of $9 \log_2(1/\varepsilon)$ for all of the above-mentioned deterministic, ancilla-free methods for general rotations.

In contrast, it was recently shown [7] that by using nondeterministic circuits that employ a small number of ancilla qubits, the number of T gates can be further reduced by a factor of 2.5 on average for axial rotations, to a count with leading term $1.15 \log_2(1/\varepsilon)$. Using Euler angles as in Eq. (1) this leads to a leading term for the expected number of T gates of $3.45 \log_2(1/\varepsilon)$ for general rotations. Remarkably, RUS circuits are able to surpass the theoretic bound for purely unitary decomposition (see Appendix A). For decomposition of a given unitary U , these so-called repeat-until-success (RUS) circuits [8] consist of repeated application of a Clifford + T sequence on an input state $|\psi\rangle$ and an ancilla qubit, followed by measurement of the ancilla qubit to project the input state $|\psi\rangle$ to the state $U|\psi\rangle$ [7,8].

An RUS circuit allows for a potentially unlimited sequence of trial and correction cycles with guaranteed finite *expected cost* below the lower bound achieved by a purely unitary circuit design. The correction circuit in each cycle can be designed to have zero cost, namely, by requiring it to be a circuit consisting only of Pauli gates. The synthesis algorithm for RUS circuits over the Clifford + T basis is based on a randomized search and achieves an expected mean T count with a leading term of $(1 + \delta) \log_2(1/\varepsilon)$, where $\delta = 0.15$ was achievable for practically important precisions ε [7].

In contrast, the probabilistic quantum circuit with fallback (PQF) protocols introduced in this work entails at most a small finite number of trials (possibly all different), and one final, purely unitary, correction step. The final correction step, or *fallback*, may have considerable cost, however, the probability of requiring the fallback step can be very small allowing for an improved expected cost for decomposition. Synthesizing a PQF circuit to approximate a given target is remarkably simpler than in the RUS case. In addition, we generalize PQF to three universal quantum bases: Clifford + T , Clifford + V [9], and Clifford + $\pi/12$. Clifford + V was previously considered for purely unitary, deterministic decomposition of single-qubit gates and resulted in the shortest known single-qubit circuits [10,11]. Clifford + $\pi/12$ has been identified as relevant for quantum computer architectures based on metaplectic anyons [12].

We present an efficient algorithm for single-qubit decomposition based on our PQF protocol. We describe the algorithmic steps in detail for each of the three bases considered. Our algorithm achieves an expected gate count of $\log_b(1/\varepsilon) + O\{\ln[\ln(1/\varepsilon)]\}$, where b is related to the scaling of the number of unique circuits that can be formed over the underlying basis. More precisely, b is defined so that for a given depth t the number of unique circuits scales as $\Theta(b^t)$, i.e., b characterizes the expansion of the underlying set of generators. Specifically, we have $b = 2$ for Clifford + T , $b = 5$ for Clifford + V , and $b = 4$ for Clifford + $\pi/12$. The PQF protocol can be generalized to several bases as the exactly representable unitaries over a given basis are representable as unitarizations of matrices over rings of cyclotomic integers of order $4m$, $m \in \{1,2,3\}$, where $m = 1$ for Clifford + V , 2 for Clifford + T , and 3 for Clifford + $\pi/12$. We aim to generalize our designs to other cyclotomic orders in future work.

II. DESIGN OF PROBABILISTIC QUANTUM CIRCUITS WITH FALLBACK

In this section, we define probabilistic quantum circuits with fallback (PQF). Our PQF protocol employs both probabilistic and deterministic subcircuits. The former are referred to as *primary* and the latter as *fallback*. We focus on the case of PQF circuits for single-qubit unitaries that are axial rotations around the Z axis, which by the Euler angle decomposition [Eq. (1)] will imply PQF protocols for arbitrary single-qubit unitaries. However, we point out that in principle the probabilistic circuit design described in this paper can also be applied to multiqubit unitaries and even to systems consisting of higher-dimensional subsystems such as, e.g., qutrits.

The primary subcircuits can be synthesized using existing synthesis methods [9,13] that given G and ε generate a probabilistic circuit $P(G,\varepsilon)$ to perform an ε approximation of the gate G with probability $p > 0$ and performs some other unitary gate G_1 with probability $1 - p$. Let $C_P(G,\varepsilon)$ be the execution cost of the $P(G,\varepsilon)$ circuit. The *fallback* subcircuit can be constructed using a synthesis method that for a given unitary target gate G and a desired precision ε generates an ε -approximation circuit $F(G,\varepsilon)$ with a known execution cost $C_F(G,\varepsilon)$, such as those in [4,6,9].

If $C_P(G,\varepsilon)$ is uniformly smaller than $pC_F(G,\varepsilon)$, then an ε approximation of the gate G using PQF will have lower

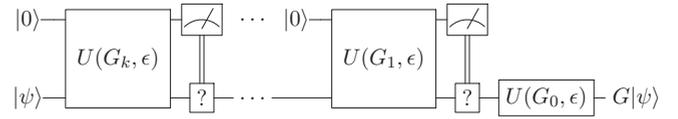


FIG. 1. PQF protocol to implement a unitary gate G .

expected cost than implementing the fallback circuit at cost $C_F(G,\varepsilon)$. To this end, we create a circuit with classical feedback that first performs the subcircuit $P(G,\varepsilon)$ with the (desired) outcome $\sim G|\psi\rangle$ upon measuring 0 or (undesired) outcome $G_1|\psi\rangle$ upon measuring 1. If the measurement outcome is 1, the circuit then performs $F(GG_1^\dagger,\varepsilon)$ on $G_1|\psi\rangle$. The expected cost of the entire circuit protocol is $C_P(G,\varepsilon) + (1 - p)C_F(GG_1^\dagger,\varepsilon)$, which is smaller than $C_F(G,\varepsilon)$ if and only if $C_P(G,\varepsilon) < pC_F(GG_1^\dagger,\varepsilon)$.

The concatenated circuit, denoted as $PQF(G,\varepsilon,1)$, is a special case of a nested probabilistic circuit, $PQF(G,\varepsilon,k)$, $k \in \mathbb{Z}$, $k \geq 0$, defined inductively as follows:

$$\begin{aligned} PQF(G,\varepsilon,0) &= F(G,\varepsilon), \\ PQF(G,\varepsilon,k) &= P(G,\varepsilon) \cup BC[PQF(GF_k^\dagger,\varepsilon,k-1)], \end{aligned} \quad (2)$$

where F_k is the undesirable outcome of the P circuit and BC denotes binary classical control on such an outcome.

The general layout of a PQF circuit is shown in Fig. 1. The “question mark” box denotes the binary classical control switch that implements the remainder of the circuit if and only if the measurement result is 1. Let $F_j|\psi\rangle$ be the undesired result upon measurement of 1 at the j th round of the protocol. Then, $G_{j-1} = G_j F_j^\dagger$ and we note that the synthesis algorithm computes $\prod_{j=0}^{k-1} G_j$.

Assuming that all rounds of the PQF circuit have a similar probability of success p and that for fixed ε the cost $C_F(G,\varepsilon)$ is similar for all target gates G , then it follows that if $PQF(G,\varepsilon,1)$ is a cost improvement over $PQF(G,\varepsilon,0)$, then $PQF(G,\varepsilon,k)$ is a cost improvement over $PQF(G,\varepsilon,k-1)$ for any $k > 0$. However, we show that the incremental improvement scales like $O[(1-p)^k]$ and therefore near-optimal performance can be achieved with a relatively small number of rounds k .

Both the PQF and RUS protocol require synthesis of unitary subcircuits. In PQF, the primary probabilistic subcircuits vary at each round, and the final fallback circuit, if necessary, is deterministic. In RUS, the same probabilistic subcircuit is applied in each round, followed by the same correction if necessary.

In Fig. 2, we compare the Markov chain [14] corresponding to the implementation of an RUS protocol to that corresponding to a PQF protocol. Both protocols implement a target single-qubit unitary transformation G for a given target approximation ε by performing a random walk on the nodes, where the target unitary G is an absorbing state, i.e., the walk terminates once it arrives in this state. In general, each node represents the unitary transformation that has been applied to the input state at the respective stage in the protocol. Both in case of RUS and PQF protocols, each transition between

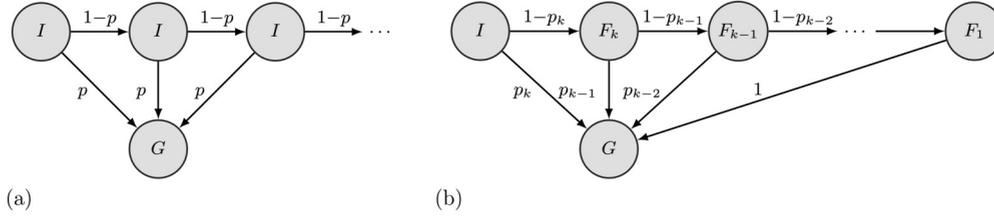


FIG. 2. Markov chains for the implementation of a target unitary transformation G . Shown are state transitions for (a) repeat-until-success (RUS) protocols and (b) probabilistic circuits with fallback (PQF) protocols.

nodes is probabilistic and is induced by the success or failure of a multiqubit unitary followed by a measurement. In the case of the RUS protocol shown in Fig. 2(a), the applied transformation in case of failure is always the identity, or, more generally, a local Clifford operation which however can easily be corrected to become the identity, whence we represent this case by the identity operator I also, whereas in case of the PQF protocol each intermediate node corresponds to an operation F_j , i.e., as shown in Fig. 1(b), the state is $F_j|\psi_j\rangle$, where F_j is the undesired result upon measurement of 1 at the j th attempt of the protocol and $|\psi_j\rangle$ was the state in the previous round, where $j = k, k - 1, \dots, 1$. The probability of success of this step is denoted by p_j and the probability of failure correspondingly by $1 - p_j$.

If we find ourselves in node j of the protocol, we not only know the entire history of previous failed attempts to implement the target gate G , we can also attempt to reach the target state G by applying a probabilistic circuit that implements $G_{j-1} := G_j F_j^\dagger$ where G_j was defined in a previous round. In any case, we will implement G after at most k steps as $G = \prod_{j=0}^{\ell} G_j$, where $\ell \in \{0, \dots, k\}$ denotes the first point in time where the protocol had a successful transition to the target gate G .

It is useful to think of the probabilistic transitions into the nodes labeled with F_j for $j = k, k - 1, \dots, 2$ as being very cheap, whereas the last transition (the “fallback”) from $F_1 \rightarrow G$ is expensive, but will always lead to the absorbing state, i.e., it guarantees that the protocol implements G with precision ε after at most k rounds.

Note that RUS designs might require a potentially unbounded number of iterations to reach the accepting state to implement the target gate G . In contrast, a PQF design with k stages is guaranteed to always implement the target gate after at most k attempts.

Clearly, in Fig. 2(a) all deterministic nodes attempt to implement one and the same target G , while in (b) the target gate of each deterministic node can and will be different. Despite this distinction, the quantum cost of a deterministic node in both Figs. 2(a) and 2(b) is defined by the set precision ε and in general this cost is about the same across all the nodes. The two main advantages of using the PQF designs is that with PQF it is easier to synthesize circuits for various different universal gate sets and that the finiteness of the designs might facilitate the layout of the circuits onto a selected fault-tolerant quantum computer architecture. The latter applies in particular if the underlying hardware supports on-demand scheduling of future operations depending on earlier measurements and if guaranteed execution time is a design objective.

III. COST ANALYSIS OF PQF CIRCUITS

The optimal T count has been proven to be an invariant of the unitary operation represented by a Clifford + T circuit [15–17]. In particular, the optimal T count is the same across various definitions of canonical and normal forms for Clifford + T circuits. At present, similar invariants have not been shown for the Clifford + $\pi/12$ basis. For the analysis that follows, the upper bounds proven in Appendix E suffice.

Consider a measurement of the ancilla qubits in the PQF design, such that one measurement outcome is labeled “favorable” and all other measurement outcomes are labeled “unfavorable.” Let the probability of the “favorable” outcome be p and the unitary applied to the target qubits upon favorable measurement be V . Let $C(U)$ be the cost of a circuit that performs U . Assuming that the cost of performing Clifford gates is negligible, the expected T count of an RUS circuit is approximately $E[C(V)] = C(U)/p$ [7].

We assume that all rounds of the PQF circuit shown in Fig. 1 have the same probability p_k of the favorable outcome that are all roughly equal to the same value, say p . Furthermore, we assume that the probability q of unfavorable outcome satisfies $q = 1 - p \ll p$ and that each round roughly has the same execution cost $C_P(\varepsilon)$. We note that all these assumptions are justified by the properties of the PQF protocol derived in the following sections and obtain the following:

Lemma 1. For a fixed ε , the following holds:

- (1) The expectation of the cost of the PQF protocol with $k > 0$ rounds is

$$C_P(\varepsilon)/p + O(q^k).$$

- (2) The variance is given by

$$C_P(\varepsilon)^2 q/p^2 + O(q^k).$$

Proof.

- (1) Let E_k be the expected cost of the k -round protocol. Assuming $q = 1 - p \ll p$ we note that $p = 1/p - q(p + 1)/p = 1/p + O(q)$. Clearly, we have that

$$\begin{aligned} E_1 &= pC_P(\varepsilon) + q[C_P(\varepsilon) + C_F(\varepsilon)] \\ &= C_P(\varepsilon)/p + O(q)C_P(\varepsilon) + q[C_P(\varepsilon) + C_F(\varepsilon)] \\ &= C_P(\varepsilon)/p + O(q). \end{aligned}$$

This provides a basis for induction on k . We have $E_{k+1} = pC_P(\varepsilon) + q[C_P(\varepsilon) + E_k] = C_P(\varepsilon) + qE_k$. By the induction hypothesis this is equal to $C_P(\varepsilon) + q[C_P(\varepsilon)/p + O(q^k)] = (p + q)C_P(\varepsilon)/p + O(q^{k+1})$ which proves claim (1).

- (2) Let $E_k^{(2)}$ be the expectation of the square of the cost. We are going to prove by induction that $E_k^{(2)} = (1 +$

$q)/p^2 C_P(\varepsilon)^2 + O(q^k)$. For the basis of the induction we observe that

$$\begin{aligned} p &= (1+q)/p^2 + p - (1+q) + O(q^2) \\ &= (1+q)/p^2 - 2q + O(q^2) \\ &= (1+q)/p^2 + O(q). \end{aligned}$$

Therefore, we obtain that

$$\begin{aligned} E_1^{(2)} &= pC_P(\varepsilon)^2 + q[C_P(\varepsilon) + C_F(\varepsilon)]^2 \\ &= (1+q)/p^2 C_P(\varepsilon)^2 + O(q)C_P(\varepsilon)^2 \\ &\quad + q[C_P(\varepsilon) + C_F(\varepsilon)]^2. \end{aligned}$$

This in turn implies

$$\begin{aligned} E_{k+1}^{(2)} &= pC_P(\varepsilon)^2 + qC_P(\varepsilon)^2 + 2qC_P(\varepsilon)E_k + qE(2)_k \\ &= C_P(\varepsilon)^2 + 2qC_P(\varepsilon)[C_P(\varepsilon)/p + O(q^k)] \\ &\quad + q[(1+q)/p^2 C_P(\varepsilon)^2 + O(q^k)] \\ &= [(p+q)^2 + q]/p^2 C_P(\varepsilon)^2 + O(q^{k+1}), \end{aligned}$$

which concludes the induction step. Thus, the variance of the cost of the k -round protocol is

$$\begin{aligned} E_k^{(2)} - E_k^2 &= (1+q)/p^2 C_P(\varepsilon)^2 \\ &\quad + O(q^k) - [C_P(\varepsilon)/p + O(q^k)]^2 \\ &= C_P(\varepsilon)^2 q/p^2 + O(q^k). \end{aligned}$$

■

Note that in contrast to unitary protocols, which are deterministic in the sense that they have a fixed cost, the cost of the PQF protocol becomes a random variable. While Lemma 1 bounds the expectation and the variance of the cost of implementing a target unitary, in a given run of the protocol we might get unlucky and require a significantly higher cost compared to the expected value. An interesting problem that ensues is to study the behavior of larger scale circuits that are composed of such probabilistic primitives. This is of particular interest in cases where the primitives are allowed to be executed in parallel and might have to wait at certain, well-defined synchronization points before a subsequent computation can proceed.

Recent work [18, Sec. 2A] considered the case of large-scale circuits composed from probabilistic elements in such a way that they can be arranged into a finite number of rounds r , each of which must complete before the next round is allowed to start. In this context, the authors established a composition theorem that asserts an exponentially decaying tail bound for the probability of the computation taking longer than $c_1 r + c_2 \ln(n)$, where n is the total number of qubits and c_1 and c_2 are constants that depend only on the maximum number of wires entering and leaving a gate and on the probability of failure. While a more detailed analysis of the exact value of these constants depends on the exact details of the implementation and is an interesting question for future research, the results of [18] imply that by making the probability of success of the gates arbitrarily close to 1, which can be achieved by choosing suitable modifiers as described in the following sections, the constant c_1 can be chosen to be arbitrarily close to 1. Hence, the overhead that arises from using probabilistic components in

larger-scale networks asymptotically leads to only a moderate, *additive* time overhead of $O[\ln(n)]$ steps. In particular, the asymptotic advantage that arises from improving the constant in front of the $\ln(1/\varepsilon)$ term carries over to larger-scale networks of probabilistic gates.

IV. OVERVIEW OF THE PQF ALGORITHM

In this section, we provide an overview of the stages of our PQF algorithm. The algorithm returns a probabilistic quantum circuit with fallback over the chosen basis that approximates a given rotation by angle θ about the z axis, denoted as $V = R_z(\theta)$, to precision ε . For a multiround PQF protocol with k rounds, the algorithm sequentially generates a subcircuit for each round. Each subsequent round of the protocol is conditional on the failure of all previous rounds, and aims to first “correct” the cumulative undesired z rotations and second to apply the target z rotation.

We develop our PQF decomposition algorithm for axial z rotations, which in turn allow the implementation of arbitrary, nonaxial rotations by Euler angle decomposition, given in Eq. (1). As a matter of principle, a PQF decomposition for a single-qubit nonaxial unitary may be directly achieved without invoking Euler angle decomposition. However, devising a practical method for such synthesis is currently an open problem.

Our algorithm takes a predefined number of PQF rounds k as input. For each round k , it generates a primary subcircuit. After determining the k subcircuits, it generates a single unitary fallback circuit that terminates the PQF protocol. The fallback circuit is only applied if all k rounds fail, and is constructed to both correct all accumulated undesired z rotations and apply the target z rotation.

The value of k can be optimized at compile time based on the appropriate cost measure. It follows from the analysis in Sec. III that the mean expected improvement in gate count of a $k+1$ -round PQF circuit over a k -round PQF circuit scales down as $O[(1-p)^k]$, where p is the typical single-round success probability. In this paper, we show that the probability can be boosted to $\Omega[1 - 1/\ln(1/\varepsilon)]$ level. Consequently, in our numerical experiments we observe insignificant improvements due to adding a second round.

The compilation stages for each round are outlined in Fig. 3. In stage 1, detailed in Sec. V, an initial approximation of the target rotation phase factor $e^{i\theta}$ is obtained. Namely, we find an algebraic number of the form z^*/z , where z belongs to a set based on the chosen gate basis, to approximate $e^{i\theta}$ by finding an approximate solution to an integer relation problem. We note that z is defined up to an arbitrary real-valued factor. The approximation is modified if needed in stage 2, by seeking either a solvable norm equation in the case of Clifford + T and Clifford + $\pi/12$ or a solvable two squares equation in the case of Clifford + V , and high success probability. Stage 2 is described in Sec. VI. A two-qubit unitary corresponding to the (modified) rational is composed in stage 3 and finally synthesized into a PQF subcircuit over the chosen basis in stage 4. If an undesired measurement outcome occurs in the current round, the undesired rotation angle ν and the next target angle $\theta - \nu$ are generated and the latter is then used, recursively, to generate the next round of the PQF protocol.

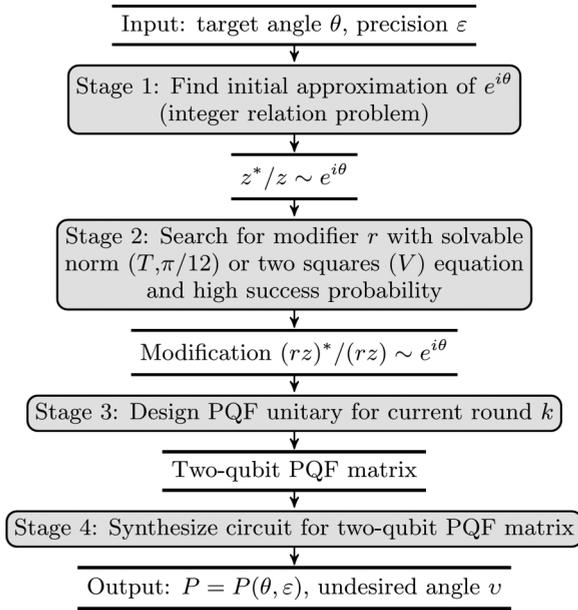


FIG. 3. Overview of compilation flow for one round of the PQF circuit. If a number of rounds k strictly larger than one is intended, stages 1–4 need to be repeated for modified target angles as described in the text.

The PQF algorithm over Clifford + T and Clifford + $\pi/12$ is detailed in Sec. VIII, and in Sec. VII for Clifford + V .

V. STAGE 1: CYCLOTOMIC RATIONAL APPROXIMATION

In this section, we review the most general stage of our synthesis method. It requires very few modifications when considering different basis sets.

Let $\zeta = e^{2\pi i/m}$ be the m th primitive root of unity and consider the corresponding ring of cyclotomic integers $\mathbb{Z}[\zeta]$. It is well known (cf. [19]) that the minimal polynomial of ζ over rationals is monic and has degree $d = \phi(m) < m$ where ϕ is the Euler totient function. We analyze the representation of an arbitrary phase factor by a unimodal cyclotomic rational z^*/z , where $z \in \mathbb{Z}[\zeta]$.

Let θ be a real angle. By direct complex expansion $|z^*/z - e^{i\theta}| = 2|\text{Im}(ze^{i\theta/2})|/|z|$. The phase factor $e^{i\theta}$ is representable exactly as z^*/z if and only if $\text{Im}(ze^{i\theta/2}) = 0$. It is approximately representable at precision ϵ if and only if $2|\text{Im}(ze^{i\theta/2})| < \epsilon|z|$. Now, consider the standard integer basis $\{1, \zeta, \dots, \zeta^{d-1}\}$ in $\mathbb{Z}[\zeta]$. Representing z in this basis results in $z = a_0 + a_1\zeta + \dots + a_{d-1}\zeta^{d-1}$, where $\{a_0, a_1, \dots, a_{d-1}\}$ are ordinary integers.

Again, by direct complex expansion we observe that $\text{Im}(ze^{i\theta/2})$ is a linear form with real coefficients in $\{a_0, a_1, \dots, a_{d-1}\}$. We expand this form as $F[a, x(\theta)] = a_0x_0(\theta) + a_1x_1(\theta) + \dots + a_{d-1}x_{d-1}(\theta)$, where $x_j(\theta) = \sin(\theta/2 + 2\pi j/m)$ and $j = 0, \dots, d-1$ is the corresponding real vector. It is easy to see that for θ in a general position, the vector does not have zero components. It is also helpful to observe that for $|\theta| < \pi/2$ at least one x_j is well separated from zero [e.g., at least one $x_j(\theta)$ has to be greater than $\sin(2\pi/m)$].

Representing the phase factor $e^{i\theta}$ exactly as a cyclotomic rational is equivalent to solving an integer relation with real coefficients, namely, solving $F[a, x(\theta)] = 0$ for a . Furthermore, when it is not solvable we consider finding approximate integer relations, i.e., finding $\{a_0, a_1, \dots, a_{d-1}\}$ that $|F[a, x(\theta)]| < \delta$. It is well known [20] that such approximate relations can be algorithmically found for arbitrarily small positive δ .

Lemma 2. For a fixed θ in a general position, $|\theta| < \pi/2$, and sufficiently small $\delta > 0$, there exists an integer solution a of $|F[a, x(\theta)]| < \delta$ such that $|a_j| = O(\delta^{-1/(d-1)})$, $j = 0, \dots, d-1$.

Proof. The proof follows from a more general theorem regarding the quality of multivariate Diophantine approximations (cf., [21], Sec. II, Theorem 1C): For any real numbers x_1, \dots, x_n and $0 < \epsilon < 1$ there exist integers q_1, \dots, q_n, p such that $|q_1x_1 + \dots + q_nx_n - p| < \epsilon$ and $\max(|q_1|, \dots, |q_n|) < \epsilon^{-1/n}$.

We apply this theorem to our case for $n = d-1$. As observed, at least one of the coefficients $x_j(\theta) = \sin(\theta/2 + 2\pi j/m)$ is in the interval $[\sin(2\pi/m), 1]$. We can relabel the $x_j(\theta)$ for convenience so that one of the coefficients belonging to the interval $[\sin(2\pi/m), 1]$ is labeled $x_0(\theta)$. Set $x_j = x_j(\theta)/x_0(\theta)$, $j = 1, \dots, d-1$. By applying Theorem 1C, we conclude that there exists an integer solution a of $|F[a, x(\theta)]|/|x_0(\theta)| < \epsilon$ with $|a_j| < \epsilon^{-1/(d-1)}$, $j = 1, \dots, d-1$.

By the triangle inequality, $|a_0| \leq |a_1||x_1| + \dots + |a_{d-1}||x_{d-1}| + \epsilon$, where ϵ is negligibly small compared to $\epsilon^{-1/(d-1)}$ and where, by design, $|x_j| < 1/|x_0(\theta)|$, $j = 1, \dots, d-1$. Thus, $|a_0| < (d-1)/|x_0(\theta)|\epsilon^{-1/(d-1)}$. Setting ϵ to be smaller than $\delta/|x_0(\theta)|$ concludes the proof of the lemma. ■

Corollary 3. For a fixed θ in a general position, $|\theta| < \pi/2$, and sufficiently small $\epsilon > 0$, there exists a cyclotomic rational approximation $|z^*/z - e^{i\theta}| < \epsilon$, $z \in \mathbb{Z}[\zeta]$ with $|z|$ in $O(\epsilon^{-1/d})$.

Proof. Per Lemma 2, a solution z to $|z^*/z - e^{i\theta}| < \epsilon$ must exist. Setting $\delta = \epsilon|z|/2$ in the lemma, we infer the existence of a solution z in $O(|z|^{-1/(d-1)}\epsilon^{-1/(d-1)})$. This implies $|z|^{d/(d-1)}$ in $O(\epsilon^{-1/(d-1)})$, or $|z|^d$ in $O(\epsilon^{-1})$, and the corollary follows. ■

In order to find the solutions algorithmically, we customize the PSLQ integer relation algorithm [20,22], where PSLQ stands for “partial sums of squares lower trapezoidal-orthogonal matrix factorization”. PSLQ is an iterative algorithm to solve integer relations of the form $a \cdot x = a_1x_1 + \dots + a_dx_d$, where $|a \cdot x|$ can be made arbitrarily small after a large enough number of iterations. Our customization terminates when the equivalent of the $|z^*/z - e^{i\theta}| < \epsilon$ inequality is first satisfied. The performance proofs in [20,22] can be modified to show that $|z|$ upon termination is in $O(\epsilon^{-1/d})$.

Our numerical experiments provide an estimate of the asymptotics of $|z|$ in cases $m = 4, 8, 12$ (which correspond to the Clifford + V , Clifford + T , and Clifford + $\pi/12$ bases, respectively). For example, when $m = 8$ (Clifford + T) we find $|z| < \kappa\epsilon^{-1/4}$, where $\kappa = 3.05 \pm 0.28$.

The following observation will be necessary for compilation stage 2 in designing a matrix over the Clifford + T and Clifford + $\pi/12$ bases. In that context, not only the size of z but also the size of its “Galois conjugate” z^* comes into play.

Observation 4. In the context of Corollary 3, consider z^\bullet which is obtained from z by formally replacing ω by $-\omega$. That is, $z^\bullet = a_0 + a_1(-\omega) + \dots + a_{d-1}(-\omega)^{d-1}$ when $z = a_0 + a_1\omega + \dots + a_{d-1}\omega^{d-1}$. The z^\bullet is also in $O(\varepsilon^{-1/d})$.

Indeed, the proof of Lemma 2 relies only on the bounds for absolute values of the coefficients, and does not change if only the signs of the coefficients are altered. Thus, z^\bullet is in $O(|z|^{-1/(d-1)}\varepsilon^{-1/(d-1)})$ and the observation follows.

Observation 5. In this context we can assume, without loss of generality, that $\ln(|z|) > 1/(2d) \ln(1/\varepsilon)$.

If this is not the case, pick the integer $s = \lceil \varepsilon^{-1/(2d)} / |z| \rceil$ and replace z with sz .

VI. STAGE 2: SEARCH FOR A MODIFIER

In this section, we present the design of our algorithm and supporting mathematical rigor required for applying the algorithm to each gate basis. Following the previous section, let $\zeta = e^{2\pi i/m}$. We limit our analysis to values of m that are multiples of 4 such that the ring $\mathbb{Z}[\zeta]$ contains $i = \sqrt{-1}$.

We introduce the *unitarization base* v , where $v = \sqrt{2}$ for $m > 4$ and $v = \sqrt{5}$ for $m = 4$ (the latter will be relevant for the V basis). Let θ be the target angle of rotation about the Z axis and z^*/z , where $z \in \mathbb{Z}[\zeta]$, be an ε approximation of the phase factor $e^{i\theta}$.

The synthesis of both purely unitary and measurement-assisted decomposition circuits hinges on the existence of a unitary matrix of the form

$$\frac{1}{v^L} \begin{bmatrix} z & y \\ -y^* & z^* \end{bmatrix}, \quad (3)$$

where $y \in \mathbb{Z}[\zeta]$ and $L \in \mathbb{Z}$. The unitary condition for this matrix $|y|^2 = v^{2L} - |z|^2$ is restrictive; the matrix need not exist for an arbitrary z as z being part of a cyclotomic rational approximation does not imply its existence.

For measurement-assisted circuit decomposition, another constraint is also relevant. Assuming the unitary matrix of the form (3) exists, we may introduce $p_1 = |z|^2/v^{2L} < 1$. For the measurement-assisted circuits to have sufficient quality we will need p_1 to be greater than $1 - O(1/L)$, which we show in the following.

For $m \neq 4$, we introduce $\rho = \zeta + \zeta^*$ and the real subring $R = \mathbb{Z}[\rho] \subset \mathbb{Z}[\zeta]$. In the special case of $m = 4$, we set $R = \mathbb{Z}$ (note that $\rho = 0$ in this special case and we would be reluctant to argue that $\mathbb{Z}[0] = \mathbb{Z}$). Any element $r \in R$ evaluates to a real number.

To address both of the above constraints in one design, we note that for any nonzero $r \in R$, we have $(rz)^*/(rz) = z^*/z$. Thus, replacing z with rz does not change the cyclotomic approximation.

Lemma 6. (Metastatement) In the above context, let $L_1 = \lceil \log_v(|z|) \rceil$. There exists an algorithmically defined subset $S_z \subset R$ of cardinality $\Theta(L_1)$ such that for any $r \in S_z$

- (1) $0 < \lceil \log_v(|rz|) \rceil - \log_v(|rz|) < O(1/L_1)$, and
- (2) $\log_v(|r|)$ is in $O[\ln(L_1)]$.

We currently do not have a proof of this lemma for arbitrary cyclotomic rings, but in the following we provide evidence that it does hold for the rings that arise from the gate sets considered in this paper.

Assuming Lemma 6 holds, for any $r \in S_z$ let $L_r = \lceil \log_v(|rz|) \rceil$. Claim 1 of Lemma 6 directly implies that $p_r = |rz|^2/v^{2L_r} > 1 - O(1/L_1)$ while claim 2 implies that L_r is in $L_1 + O[\ln(L_1)]$. There is also an algorithmically defined set of at least $\Theta(L_1)$ values with these properties.

Intuitively, in the subsequent designs for the probabilistic measurement-assisted circuits $p_r > 1 - O(1/L_1)$ means there will be a high one-round success rate. The asymptotics for L_r implies that no r from S_z will substantially increase the depth of the resulting circuit.

The existence of the matrix

$$\frac{1}{v^{L_r}} \begin{bmatrix} rz & y \\ -y^* & rz^* \end{bmatrix}, \quad (4)$$

for some chosen $r \in S_z$, is equivalent to solving the equation

$$|y|^2 = v^{2L_r} - |rz|^2 \quad (5)$$

for $y \in \mathbb{Z}[\zeta]$. Equation (5) is known as a *norm equation* over the cyclotomic integers. Its solvability and solutions are well understood [19].

Consider the *absolute norm* map $N : \mathbb{Q}(\rho) \rightarrow \mathbb{Q}$. It is a general fact that $N(R) \subset \mathbb{Z}$. We use terminology and facts from [19] in our following description. We first address a particular case where $p = N(v^{2L_r} - |rz|^2)$ is a prime integer. As per Theorem 2.13 [19], the norm equation in Eq. (5) is solvable if and only if $p \equiv 1 \pmod{m}$. Intuitively, this means that solvable norm equations are not rare. Let B be an arbitrarily large positive integer. It is well known that the density of prime numbers in, say, the segment $[B/2, B]$ is in $\Omega[1/\ln(B)]$. It is also well known that if $m \ll B$, then the density of such prime numbers p such that $p \equiv 1 \pmod{m}$ in that segment is still in $\Omega[1/\ln(B)]$.

Suppose we have identified a large enough subset $S_z \subset R$ so that the set of integers $\{N(v^{2L_r} - |rz|^2) | r \in S_z\}$ intersects with some segment of the form $[B/2, B]$ and the intersection has $\Theta[\ln(B)]$ distinct integers, i.e., the number of distinct integers in the intersection is $\ln(B)$ times some significant factor. Then, with some high probability there is an $r \in S_z$ such that $p = N(v^{2L_r} - |rz|^2)$ is prime and $p \equiv 1 \pmod{m}$.

However, this is a minimalistic approach. If $p = N(v^{2L_r} - |rz|^2)$ is not prime, and its prime factorization is known, then the complete analysis of solvability of Eq. (5) can be algorithmically performed in polynomial time. Therefore, we can broaden the search for feasible values of r by looking at such values where $N(v^{2L_r} - |rz|^2)$ is easy to factor (e.g., it is a smooth integer). We refer to Eq. (5) as *easily solvable* when $N(v^{2L_r} - |rz|^2)$ is easy to factor *and* the equation has a solution.

Conjecture 7. (Metaconjecture) For any $z \in \mathbb{Z}[\zeta]$ in a general position there exists a certain subset $S_z \subset R$ that satisfies the claims of Lemma 6, has cardinality in $O(L_1)$, and contains at least one r for which Eq. (5) is easily solvable.

Assuming this conjecture holds, we can manufacture a unitary matrix of the form (4) in polynomial classical runtime with the promise that its L_r is in $\log_v(|z|) + O\{\ln[\ln(|z|)]\}$. We proceed by describing how to apply the described framework in the context of each of the three basis sets.

VII. PQF OVER CLIFFORD + V

Unitary decomposition of single-qubit rotations over the so-called V basis was described in Ref. [9]. While RUS decomposition over Clifford + V has not yet been shown, the PQF protocol allows generalization to the V basis. We generalize PQF to Clifford + V in this section, and show that it is remarkably simpler than in the Clifford + T case.

Recall that the single-qubit V gate is given by $V = (I - 2iZ)/\sqrt{5}$. The group of circuits generated by the Clifford group and the V gate is universal for quantum computation [9]. It has also been shown that an arbitrary single-qubit unitary gate can be approximated to precision ϵ by a single-qubit unitary Clifford + V circuit with V count bounded by $3 \log_5(1/\epsilon) + O\{\ln\{\ln(1/\epsilon)\}\}$ [10]. For approximation of axial rotations, the algorithm is efficient. For arbitrary single-qubit targets, the same V count can be achieved using an exponential-time algorithm [9] that is practically feasible for a reasonable range of precisions.

The guarantees for the unitary decomposition algorithm are based on the following result which was shown in [9] and which is the basis for our efficient algorithm to approximate an arbitrary axial rotation $R_z(\theta)$ with a PQF circuit over Clifford + V :

Lemma 8. A unitary matrix of the form $\frac{1}{\sqrt{5}^L} \begin{bmatrix} z & y \\ -y^* & z^* \end{bmatrix}$, where y, z are Gaussian integers and $L \in \mathbb{Z}$, can be exactly and algorithmically decomposed into a Clifford + V circuit of V count at most L .

Following Lemma 8, we consider the case $\zeta = i$ (which corresponds to $m = 4$). Here, $\mathbb{Z}[i]$ is a quadratic extension of \mathbb{Z} so $d = 2$. By Corollary 3, any phase factor $e^{i\theta}$ can be approximated with a Gaussian rational $z^*/z, z \in \mathbb{Z}[i]$, where $|z|$ is in $O(\epsilon^{-1/2})$.

Observation 9. Without loss of generality, we can assume that $\ln(|z|) > 1/4 \ln(1/\epsilon)$.

If it is not the case, we pick the integer $s = \lceil \epsilon^{-1/4}/|z| \rceil$ and replace z with $s z$.

A. Stage 2: Probability modifier

We now prove Lemma 6 over Clifford + V . Let $L_1 = \lceil \log_{\sqrt{5}}(|z|) \rceil$. We want to define a subset $S_z \subset \mathbb{Z}$ such that $\forall r \in S_z, \lceil \log_{\sqrt{5}}(|rz|) \rceil - \log_{\sqrt{5}}(|rz|) < 1/L_1$. Let $\lambda = L_1 - \log_{\sqrt{5}}(|z|)$. We select values of r such that $0 < \log_{\sqrt{5}}(|r|) - \lfloor \log_{\sqrt{5}}(|r|) \rfloor < \lambda$. Under this assumption, $\lceil \log_{\sqrt{5}}(|rz|) \rceil - \log_{\sqrt{5}}(|rz|) = \lambda - (\log_{\sqrt{5}}(|r|) - \lfloor \log_{\sqrt{5}}(|r|) \rfloor)$.

We define the desired S_z as a subset of positive integers r satisfying the inequality

$$\lambda - 1/L_1 < \log_{\sqrt{5}}(|r|) - \lfloor \log_{\sqrt{5}}(|r|) \rfloor < \lambda.$$

It is necessary and sufficient that $\log_{\sqrt{5}}(|r|)$ is in an interval of the form $(k + \lambda - 1/L_1, k + \lambda)$, where $k \in \mathbb{Z}$ or $|r|$ is in the interval $I_k = (\sqrt{5}^k \sqrt{5}^{\lambda-1/L_1}, \sqrt{5}^k \sqrt{5}^\lambda)$. It follows that the number of integers contained in I_k grows exponentially with $k \geq k_0 = \lceil \log_{\sqrt{5}}(L_1) \rceil$.

We define the desired S_z as the ordered sequence of all integers in $\cup_{k \geq k_0} I_k$. While S_z is an infinite sequence, any of its initial subsequences of length $O(L_1)$ is contained in $\cup_{k_0 \leq k \leq O\{\ln(L_1)\}} I_k$. This concludes the proof.

Next, we specialize Conjecture 7 for $m = 4$ and discuss its implications. Let S'_z be some initial subsequence of length $\Theta(L_1)$ in S_z and let $S'_z(\text{easy}) \subset S'_z$ be the subsequence of such $r \in S'_z$ for which the norm equation $|y|^2 = 5_r^L - |rz|^2$ is easily solvable for $y \in \mathbb{Z}[i]$.

Conjecture 10. The density of $S'_z(\text{easy})$ in S'_z is in $\Omega[1/\ln(L_1)]$. It suffices to inspect $O(L_1)$ initial values in S_z in order to find one for which the norm equation is easily solvable.

Conjecture 10 implies that we need to test at most $O(L_1) = O\{\ln(|z|)\} = O\{\ln(1/\epsilon)\}$ norm equations for easy solvability to find one that is easily solvable. It also implies that the value of L_r corresponding to the solution is in $L_1 + O\{\ln(L_1)\} = \log_{\sqrt{5}}(|z|) + O\{\ln\{\ln(|z|)\}\} = \log_5(1/\epsilon) + O\{\ln\{\ln(1/\epsilon)\}\}$.

B. Stages 3, 4: Design and synthesis of PQF subcircuits

We now have the unitary matrix $W = \frac{1}{\sqrt{5}^{L_r}} \begin{bmatrix} rz & y \\ -y^* & rz^* \end{bmatrix}$, where $L_r = \lceil \log_5(r^2|z|^2) \rceil \leq \log_5(|z|^2) + O\{\ln\{\ln(|z|^2)\}\}$ and $r^2|z|^2/5^{L_r} > 1 - 1/\log_5(|z|^2)$.

As observed in [9], V is exactly represented by a unitary Clifford + V circuit with V count at most L_r . Therefore, the two-qubit PQF matrix $U = \text{CNOT}(I \otimes W)\text{CNOT}$ is exactly represented by a circuit with the same V count.

By direct computation, when U is applied to $|\psi\rangle|0\rangle$ and the second qubit is measured, then either

- (i) on measurement outcome 0 the $\Lambda(z^*/z) \sim \Lambda(e^{i\theta})$ rotation gate is effectively applied to the primary qubit, or
- (ii) on measurement outcome 1 the $\Lambda(-y/y^*)$ rotation gate is applied to the primary qubit.

Thus, the round-one fallback circuit must be a unitary ϵ approximation of the rotation gate $\Lambda(-y^*/ye^{i\theta})$. Fallback circuits at subsequent rounds have similar structure.

As per [10], any of the fallback circuits can be implemented at V count of at most $3 \log_5(1/\epsilon) + O\{\ln\{\ln(1/\epsilon)\}\}$. Following the guarantees derived for stage 1 of this algorithm, $|z|^2$ is in $O(1/\epsilon)$, therefore, the V count of the two-qubit circuit for $U = \text{CNOT}(I \otimes V)\text{CNOT}$ is bounded by $\log_5(1/\epsilon) + O\{\ln\{\ln(1/\epsilon)\}\}$.

The one round “failure” rate of the circuit, which is the probability q of measuring 1, is less than $1/\log_5(|z|^2)$. As per Observation 9, we can assume that $\log_5(|z|^2) > 1/2 \log_5(1/\epsilon)$ and therefore $q < 2/\log_5(1/\epsilon)$. Thus, the expected V count of the one-round PQF protocol is bounded by $\log_5(1/\epsilon) + O\{\ln\{\ln(1/\epsilon)\}\} + q(3 \log_5(1/\epsilon) + O\{\ln\{\ln(1/\epsilon)\}\}) < \log_5(1/\epsilon) + O\{\ln\{\ln(1/\epsilon)\}\}$. Similarly, the V count of the two-round PQF protocol is bounded by $\log_5(1/\epsilon) + O\{\ln\{\ln(1/\epsilon)\}\}$.

VIII. PQF OVER CLIFFORD + T AND CLIFFORD + $\pi/12$

The Clifford + T basis is arguably the most popular universal quantum basis [1]. It consists of the multiqubit Clifford group and the single-qubit T gate, where $T = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{bmatrix}$. Alternatively, the basis can be viewed as being generated by $\{T, H, \text{CNOT}\}$. We cost our synthesized circuits by the number of T gates, which is motivated by the high cost of fault-tolerant implementations of the T gate (or other non-Clifford gate) [23–25].

The Clifford $+\pi/12$ basis analogously consists of the multiqubit Clifford group and the single-qubit $K = \pi/12$ gate, where $K = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/6} \end{bmatrix}$. It is generated by the set $\{K, H, \text{CNOT}\}$. The study of this set is motivated by recent results on the universality of metaplectic anyons [12]. While we present an algorithm to decompose into this basis, an efficient fault-tolerant implementation of the $K = \pi/12$ gate, for example by magic state distillation, remains open for future research. In this analysis, however, we assume, paralleling the Clifford $+T$ basis, that the cost of executing a $\pi/12$ gate is significantly higher than the cost of executing a Clifford gate. We also assume that K, K^{-1}, K^2 , and K^{-2} have the same unit cost. Therefore, the cost of a Clifford $+\pi/12$ circuit is dominated by the number of K monomials occurring in the circuit.

In all other technical aspects, the Clifford $+T$ and Clifford $+\pi/12$ systems are strikingly similar. For the Clifford $+T$ and Clifford $+\pi/12$ bases, the first stage of the algorithm approximates the phase factor $e^{i\theta}$ with a unimodal cyclotomic rational, i.e., an algebraic number of the form z^*/z , where $z \in \mathbb{Z}[\omega]$, by finding an approximate solution of an integer relation problem. The second stage performs the modification $z \mapsto (rz)$, where $r \in \mathbb{Z}[\rho]$ using Lemma 12 developed below. The third and fourth stages design and synthesize the PQF subcircuit. We review the algorithm in the case of these two bases in the following.

A. Exactly representable unitaries

We denote $\omega = e^{i\pi/4}$ and $\omega_{12} = e^{i\pi/6}$. Both ω and ω_{12} are algebraic integers of degree 4. We intentionally omit the superscript 12 and use ω when the algebra of these two algebraic integers is identical, and denote it otherwise. We denote $\rho = \omega + \omega^* = \sqrt{2}$ and $\rho_{12} = \omega_{12} + \omega_{12}^* = \sqrt{3}$. Again, we use ρ without subscript when no distinction is necessary. The fundamental unit $\nu = 1 + \rho$ of the $\mathbb{Z}[\rho]$ ring and the fundamental unit $\nu_{12} = 2 + \rho_{12}$ of the $\mathbb{Z}[\rho_{12}]$ ring. We use ν without subscript when distinction is unnecessary.

The algebraic number ring $\mathbb{Z}[\omega]$ is a degree 4 extension of \mathbb{Z} . The Galois group of this extension is the direct product $\mathbb{Z}_2 \times \mathbb{Z}_2$ generated by complex conjugation $*$ and one other automorphism \bullet that extends $\omega^\bullet = -\omega$. The ring $\mathbb{Z}[\omega]$ has an integer basis of four elements, with the most obvious basis being $\{\omega^3, \omega^2, \omega, 1\}$ [13]. It consists of all numbers of the form $a\omega^3 + b\omega^2 + c\omega + d$, where a, b, c, d are arbitrary integers.

It was shown in [13] for the Clifford $+T$ system that a unitary operation V on n qubits is representable exactly by a Clifford $+T$ circuit if and only if it is of the form $V = 1/\sqrt{2}^k M$, where M is a matrix with elements from $\mathbb{Z}[\omega]$ and k is some non-negative integer. To satisfy the unitary condition, we require $MM^\dagger = 2^k \mathbf{1}_{2^n}$. Moreover, it was shown that a matrix of this form can be represented as an asymptotically optimal Clifford $+T$ circuit using at most two ancilla qubits [6,26], and no ancilla qubits when either the target is a single-qubit unitary or when $\det(1/\sqrt{2}^k M) = 1$ [26].

In Appendix C, we extend this claim to single-qubit unitaries over the Clifford $+\pi/12$ basis. We prove that a $V \in U(2)$ is representable exactly as a Clifford $+\pi/12$ circuit if it is of the form $V = 1/\sqrt{2}^k M$, where M is a 2×2 matrix

over $\mathbb{Z}[\omega_{12}]$ such that $MM^\dagger = 2^k \mathbf{1}_2$. Additional details on the norm equation in this case are given in Appendix B.

We note that the PQF and RUS protocols share that when the phase factor $e^{i\theta}$ is approximated by some $y/\sqrt{2}^L$, where $y \in \mathbb{Z}[\omega]$, and by some z^*/z , where $z \in \mathbb{Z}[\omega]$ for the same precision, then z will in general have much smaller bit size than y for that precision.

B. Stage 2: Probability modifier

Let $z^*/z, z \in \mathbb{Z}[\omega]$ be a cyclotomic rational approximation of $e^{i\theta}$ as explained in Sec. V. In stage 2, we include z in a unitary of the form (3), where in this context $\nu = \sqrt{2}$, $y \in \mathbb{Z}[\omega]$, and $L \in \mathbb{Z}$. We would like $|z|^2/2^L$ to be reasonably large since this value equals the success probability of the current round in the PQF protocol. Unfortunately, the majority of z values do not allow this. To create a unitary of the form (3), we seek a y that satisfies the normalization condition $(|y|^2 + |z|^2)/2^L = 1$ or, equivalently, $|y|^2 = 2^L - |z|^2$. It is easy to see that $|z|^2$ belongs to the real-valued ring $\mathbb{Z}[\rho]$ and thus so does $2^L - |z|^2$.

Given an arbitrary $\xi \in \mathbb{Z}[\rho]$, the identity

$$|y|^2 = \xi, \quad (6)$$

considered as an equation for an unknown $y \in \mathbb{Z}[\omega]$, is called a *norm equation in $\mathbb{Z}[\omega]$* . Deciding whether a given norm equation is solvable and finding a solution is in general at least as hard as performing factorization of an arbitrary integer. For our algorithm to be efficient, we need to find norm equations that are easy to solve.

A necessary solvability condition to construct a matrix of the form (3) with $\nu = \sqrt{2}$ is given by $|z|^2 \leq 2^L$ and $|z^\bullet|^2 \leq 2^L$, where $()^\bullet : \mathbb{Z}[\omega] \rightarrow \mathbb{Z}[\omega]$ extends the map $\omega \mapsto (-\omega)$. We can arbitrarily replace z in z^*/z by rz , where $r \in \mathbb{Z}[\rho]$, without changing the fraction.

Our strategy generalizes that of Ref. [3]. Consider a fixed $z \in \mathbb{Z}[\omega]$. Introduce $L_1 = \lceil \log_2(|z|^2) \rceil$. Note that for z defined at stage 1, L_1 follows the asymptotics of $\log_2(1/\varepsilon)/2 + c$, where c is a constant. Thus, for asymptotically small ε , L_1 is asymptotically large. For a randomly chosen $r \in \mathbb{Z}[\rho]$, we set $L_r = \lceil \log_2(|rz|^2) \rceil$.

Further design and analysis of our algorithm is based on the following:

Conjecture 11. In the above context, consider the set $S_+(z) = \{r \in \mathbb{Z}[\rho] \mid 2^{L_r} - |rz|^2 \geq 0, 2^{L_r} - |(rz)^\bullet|^2 \geq 0\}$. Let $S_{\text{solvable}}(z) \subset S_+(z)$ be the subset of such $r \in \mathbb{Z}[\rho]$ for which the norm equation $|y|^2 = 2^{L_r} - |rz|^2$ is solvable in $\mathbb{Z}[\omega]$. Then, the average density of $S_{\text{solvable}}(z)$ in $S_+(z)$ belongs to $\Omega(1/L_1)$, when $|z| \rightarrow \infty$ and $L_1 \rightarrow \infty$.

This conjecture is a special case of Conjecture 7. The motivation for conjectures of this type was discussed in Sec. VI. In our numeric experiments for over 1000 random angles and 30 levels of precision we have not encountered a single failure, suggesting the conjecture does hold.

Lemma 12. For sufficiently large L_1 , an $r \in \mathbb{Z}[\rho]$ can be algorithmically found, in a probabilistically polynomial

number of steps, such that

- (1) $L_r \leq L_1 + \log_2(|z^\bullet/z|) + O[\ln(L_1)]$,
- (2) $p(r) > |rz|^2/2^{L_r} > 1 - 1/L_1$,
- (3) $2^{L_r} - |r^\bullet z^\bullet|^2 \geq 0$,
- (4) the norm equation $|y|^2 = 2^{L_r} - |rz|^2$ is easily solvable in $\mathbb{Z}[\omega]$.

Proof. As per Conjecture 11, it suffices to find a set of values r that satisfy (1)–(3) in Lemma 12 of size $O(L_1)$. Then, there is at least one value r in a set of size $O(L_1)$ for which the norm equation is easily solvable. Before algorithmically constructing such a set, we note that condition 3 is somewhat redundant. It is a necessary condition for the norm equation to be solvable and conveniently helps reduce the search space for r .

Let $\zeta = \log_2(2^{L_1}/|z|^2) = L_1 - \log_2(|z|^2)$. We rewrite conditions 1 and 2 in terms of $R = \lceil \log_2(|r|^2) \rceil$ and $f = \log_2(|r|^2) - R$. The first condition can be restated as R is in $\log_2(|z^\bullet/z|) + O[\ln(L_1)]$. The second condition means that $2^{-(\zeta-f)} > 1 - 1/L_1$. We subsequently design f to be smaller than, but close enough to, ζ . It follows that $(1 - 1/L_1)2^\zeta < 2^f < 2^\zeta$. As per the definitions of ζ, f, R a slightly stronger (asymptotically equivalent) condition on $|r|$ is given by

$$[1 - 1/(2L_1)]2^{(R+\zeta)/2} < |r| < 2^{(R+\zeta)/2}.$$

We then rewrite condition 3 as $|r^\bullet| \leq 2^{L_1/2}/|z^\bullet| \times 2^{R/2}$.

We are now ready to describe the construction of a sufficient set of values r . This is one of the few places where the distinction between Clifford + T and Clifford + $\pi/12$ must be made. Recall the fundamental unit $v = 1 + \rho$ we have defined for Clifford + T and the fundamental unit $v_{12} = 2 + \rho_{12}$ defined for Clifford + $\pi/12$.

Our construction for Clifford + T exploits Lemma 17 from [3]: given real numbers x_0, x_1, y_0, y_1 such that $|(x_1 - x_0)(y_1 - y_0)| > v^2$, one can algorithmically find an $r \in \mathbb{Z}[\sqrt{2} = \rho]$ such that $r \in (x_0, x_1)$ and $r^\bullet \in (y_0, y_1)$.

The construction for Clifford + $\pi/12$ is based on a version of the lemma developed in Appendix D: given real numbers x_0, x_1, y_0, y_1 such that $|(x_1 - x_0)(y_1 - y_0)| > v_{12}^2$, one can algorithmically find an $r \in \mathbb{Z}[\sqrt{3} = \rho_{12}]$ such that $r \in (x_0, x_1)$ and $r^\bullet \in (y_0, y_1)$.

The two lemmas are identical except for the 12 subscript. We omit the subscript in the following narration which is common for the two cases. We collectively refer to the two lemmas as the ‘‘bullet lemmas.’’ Set $x_0(R) = [1 - 1/(2L_1)]2^{(R+\zeta)/2}$, $x_1(R) = 2^{(R+\zeta)/2}$, $y_0(R) = -2^{L_1/2}/|z^\bullet| \times 2^{R/2}$, and $y_1 = +2^{L_1/2}/|z^\bullet| \times 2^{R/2}$. Then, $|[x_1(R) - x_0(R)][y_1(R) - y_0(R)]| = 2^R |z/z^\bullet|/L_1$. If the latter value is greater than v^2 or, equivalently, $2^R > v^2 |z^\bullet/z| L_1 \in O(L_1)$, then one can algorithmically find at least one $r \in \mathbb{Z}[\rho]$ that satisfies conditions 1–3. Consider $R_0(z) = \lceil \log_2(v^2 |z^\bullet/z| L_1) \rceil$. Obviously, $R_0(z)$ is in $\log_2(|z^\bullet/z| L_1) + O[\ln(L_1)] = \log_2(|z^\bullet/z|) + O[\ln(L_1)]$. We note that $|[x_1(R) - x_0(R)]|$ grows exponentially with $R \geq R_0$.

Having defined $S_z(R)$ as an ordered sequence of $r \in \mathbb{Z}[\rho]$ such that $r \in [x_0(R'), x_1(R')]$ for some $R' \in [R_0, R]$, we conclude that the cardinality of $S_z(R)$ grows exponentially with $R \geq R_0$. Indeed, when $|[x_1(R) - x_0(R)]|$ is exponentially large, we can subdivide it into exponentially large number of subsegments, each minimally satisfying the condition of

Input: $z \in \mathbb{Z}[\omega]$ ▷ hyperparameters ρ, v
1: **procedure** NORMALIZATION-1(z)
2: $L_1 \leftarrow \lceil \log_2(|z|^2) \rceil, \zeta \leftarrow L_1 - \log_2(|z|^2), Y \leftarrow \text{None}$
3: $R_0 \leftarrow \lceil \log_2(|z^\bullet|/|z| L_1 v^2) \rceil$
4: $x_{min} \leftarrow (1 - 1/(2L_1))2^{(R_0+\zeta)/2}, x_{max} \leftarrow 2^{(R_0+\zeta)/2}$
5: $\Delta \leftarrow x_{max} - x_{min}$
6: **while** $Y = \text{None}$ **do**
7: $x_0 \leftarrow x_{min}, x_1 \leftarrow x_0 + \Delta$
8: **while** $Y = \text{None}$ and $x_1 \leq x_{max}$ **do**
9: construct $r \in \mathbb{Z}[\rho], r \in (x_0, x_1)$ such that
10: $|r^\bullet| \leq 2^{(R_0+\zeta)/2}$
11: $L_r \leftarrow \lceil \log_2(|rz|^2) \rceil$
12: **if** $|y|^2 = 2^{L_r} - |rz|^2$ is easily solvable **then**
13: $Y \leftarrow \{r, y\}$
14: **end if**
15: $x_0 \leftarrow x_1, x_1 \leftarrow x_1 + \Delta$
16: **end while**
17: $x_{min} \leftarrow 2x_{min}, x_{max} \leftarrow 2x_{max}$
18: **end while**
19: **end procedure**
Output: Y ▷ acceptable norm equation solution

FIG. 4. Algorithm to find a probability modifier r .

the appropriate bullet lemma, and algorithmically find an element of the $\mathbb{Z}[\rho]$ in each subsegment. Therefore, if an initial subsequence of any $S_z(R)$ has cardinality in $O(L_1)$, then that subsequence is also contained in some $S_z(R')$, where R' is in $\log_2(|z^\bullet/z|) + O[\ln(L_1)]$.

Per Conjecture 11, it would be sufficient to inspect $O(L_1)$ initial elements of a large enough $S_z(R)$ in order to find $r \in S_z(R)$ such that the norm equation $|y|^2 = 2^{L_r} - |rz|^2$ is solvable. By the above observation, such an r will be also in some $S_z(R')$, where $R' = \log_2(|z^\bullet/z|) + O[\ln(L_1)]$, and the lemma follows. ■

Corollary 13. In the context of Lemma 12 one can efficiently find an $r \in \mathbb{Z}[\rho]$ such that

$$L_r \leq \log_2(1/\varepsilon)/2 + O\{\ln[\ln(1/\varepsilon)]\} + c,$$

where c is a constant.

Proof. We find an $r \in \mathbb{Z}[\rho]$ that satisfies the conditions of Lemma 12. Recall that up to an additive fractional part $L_1 = 2 \log_2(|z|)$. By condition 1, $L_r \leq \log_2(|z|) + \log_2(|z^\bullet|) + O[\ln[\ln(|z|)]]$. As per Corollary 3 and Observation 4, both $|z|$ and $|z^\bullet|$ are in $O(\varepsilon^{-1/4})$ and our claim follows. ■

We now have an algorithm for stage 2 that iterates through a sufficient set of candidate values of r until one yields a solvable norm equation. The pseudocode is shown in Fig. 4.

C. Stage 3: PQF unitary design

When the algorithm to modify the probability succeeds for a given z , we can construct a single-qubit unitary V of the form (3), where $y, z \in \mathbb{Z}[\omega]$, $L \in \mathbb{Z}$, $v = \sqrt{2}$, and the probability of success of the current round is $|z|^2/2^L > 1 - 1/L$. For Clifford + T , the unitary V can be decomposed exactly into an optimal ancilla-free Clifford + T circuit using methods in [13]. For Clifford + $\pi/12$, we decompose it using a similar technique described in Appendix C.

The following theorem summarizes the theoretical upper bounds on the mean expected cost of a PQF circuit over the Clifford + T and Clifford + $\pi/12$ bases. For completeness,

we highlight that the same bound in fact applies to RUS circuits over the Clifford + T , resulting in a small but definitive asymptotic improvement over the bound given in Ref. [7].

Theorem 14. In the context of both PQF and RUS protocols where modifier sampling is based on Lemma 12, note the following:

(1) For the Clifford + T basis, if the T cost of the fallback round of PQF is in $O\{\ln(1/\varepsilon)\}$ then the overall expected T cost of a one-round PQF protocol is

$$\log_2(1/\varepsilon) + O\{\ln\{\ln(1/\varepsilon)\}\}. \quad (7)$$

(2) The expected T cost of an RUS protocol is also given by Eq. (7).

(3) For the Clifford + $\pi/12$ basis, the overall expected K cost of a one-round PQF protocol is

$$1/2 \log_2(1/\varepsilon) + O\{\ln\{\ln(1/\varepsilon)\}\}. \quad (8)$$

Proof.

(1) For the PQF protocol over the Clifford + T basis, the expected T cost is $2L_r + C_F[1 - p(r)]$, where C_F is the fallback cost. As per the above and Observation 5, $1 - p(r) < 2/\log_2(1/\varepsilon)$ and by Corollary 13, the claim follows.

(2) For the RUS protocol with Lemma 12, the expected T cost is $2L_r/p(r)$. As per condition 2 of the lemma, the expected cost is dominated by $2L_r(1 + 1/L_1) = 2L_r[1 + 2/\log_2(1/\varepsilon)]$ and the claim follows from Corollary 13.

(3) For the PQF protocol over the Clifford + $\pi/12$ basis, the expected K cost is bounded by $L_r + 2 + C_F[1 - p(r)]$, where C_F is the fallback cost. As per the above and Observation 5, $1 - p(r) < 2/\log_2(1/\varepsilon)$ and by Corollary 13 the claim follows. ■

D. Stage 4: Synthesis of PQF subcircuit

From the unitary matrix V , we construct a two-qubit unitary U given by

$$U = \text{CNOT}(I \otimes V)\text{CNOT} = \begin{bmatrix} V & 0 \\ 0 & XVX \end{bmatrix}.$$

We denote the primary input state for round k as $|\psi_k\rangle$. The subcircuit U for round k acts on the state $|\psi_k\rangle \otimes |0\rangle$, where the second qubit is an ancilla. We then measure the second (ancilla) qubit. When the measurement outcome is 0, the first qubit is left in the state $\begin{bmatrix} 1 & 0 \\ 0 & z^*/z \end{bmatrix} |\psi_k\rangle$ which is the desired ε approximation of $R_z(\theta)$. When the measurement outcome is 1, the first qubit is left in the state $\begin{bmatrix} 1 & 0 \\ 0 & -y/y^* \end{bmatrix} |\psi_k\rangle$. Unless $-y/y^*$ is ε close to $e^{i\theta}$, in this case we must apply the rotation $R_z(\theta')$, where $\theta' = \theta - \arg(-y/y^*)$ in the next round.

The unitary U at round k has the same T count (K count) as the T count (K count) of the optimal single-qubit Clifford + T (Clifford + $\pi/12$) circuit for unitary V since we invoke the optimal single-qubit deterministic decomposition of V to obtain its circuit. The only other gates involved are two CNOT gates. For Clifford + T , the techniques in Refs. [3–6] can be used to optimally decompose V . In Appendices C–E, we show how to optimally decompose a single-qubit gate into the Clifford + $\pi/12$ basis. The T count (K count) of the two-qubit unitary at any subsequent round is defined (asymptotically) by

Expected T-count for PQF over Clifford+T

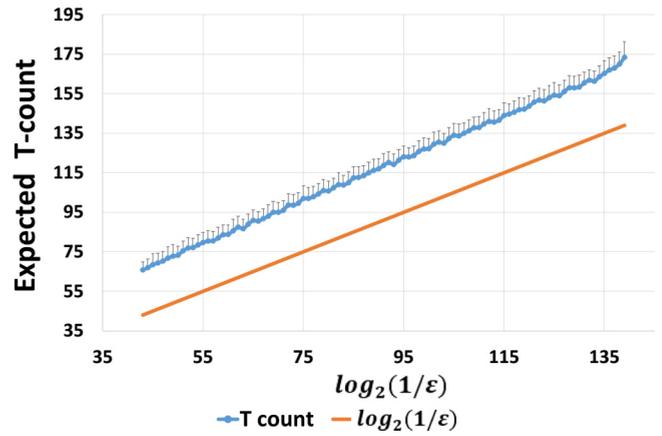


FIG. 5. (Color online) Precision ε versus mean expected T count of PQF circuits over Clifford + T for set of random angles.

the precision ε . The difference in cost between the rounds is asymptotically bounded by an $O\{\ln\{\ln(1/\varepsilon)\}\}$ term.

IX. NUMERICAL RESULTS

We evaluate the performance of our algorithm on a set of 1000 angles randomly drawn from the interval $(0, \pi/2)$ at 30 target precisions $\varepsilon \in \{10^{-11}, \dots, 10^{-40}\}$. In all numerical experiments, expected cost statistics have been collected for one-round PQF circuits. Adding the second round to the compiled circuits only improves the mean expected gate count by three gates on average. This is due to the probability modification during stage 2 for each round k of PQF compilation. Modification boosts the probability of success to typical values above 0.97 and above 0.985 for at least half of the cases.

Figure 5 plots the precision ε versus the mean (and standard deviation) of the expected T count across the PQF circuits generated for the set of 1000 random angles. The maximum likelihood estimate for the mean expected T count is $\log_2(1/\varepsilon) + 4 \log_2[\log_2(1/\varepsilon)] + 1.187$.

Expected K-count for Clifford+pi/12

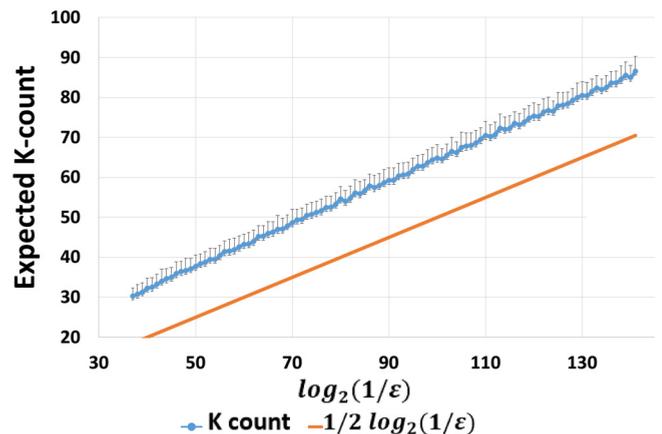


FIG. 6. (Color online) Precision ε versus mean expected K count of PQF circuits over Clifford + $\frac{\pi}{12}$ for set of random angles.

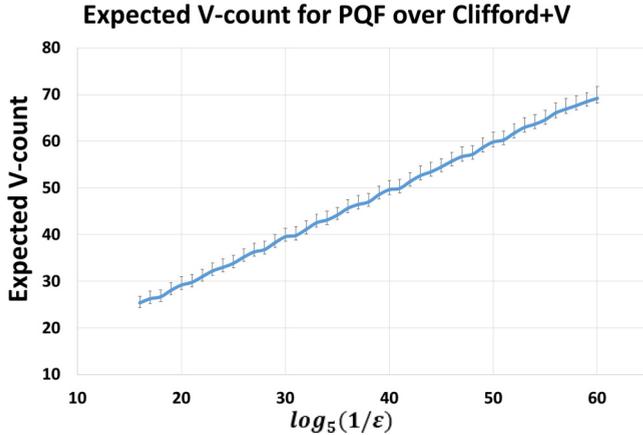


FIG. 7. (Color online) Precision ϵ versus mean expected V count of PQF circuits over Clifford + V for set of random angles.

Figure 6 plots the precision ϵ versus the mean (and standard deviation) of the expected K count for the PQF Clifford + $\pi/12$ circuits generated for the set of 1000 random angles. The maximum likelihood estimate for the mean expected K count is $1/2 \log_2(1/\epsilon) + 2 \log_2[\log_2(1/\epsilon)] + 3.48$.

Figure 7 plots the precision ϵ versus the mean (and standard deviation) of the expected V count for the PQF Clifford + V circuits generated for the set of 1000 random angles. The maximum likelihood estimate for the mean expected V count is $\log_5(1/\epsilon) + 0.95 \log_5[\log_5(1/\epsilon)] + 7.26$.

X. CONCLUSION AND FUTURE WORK

We have established a method of synthesizing circuits of certain type called probabilistic quantum circuits with fallback (PQF), which is more general and less technically challenging than the synthesis of repeat-until-success (RUS) circuits. We have demonstrated that the method can be applied to implementing single-qubit unitaries over at least three different universal quantum bases. The mean expected cost of resulting probabilistic circuits has an asymptotic upper bound with a leading term three times smaller than the leading term of the corresponding optimal purely unitary ancilla-free circuit over the same universal quantum circuit.

The design and cost analysis of the PQF circuits is performed under conjectures that are substantially the same as the norm density conjectures used in [3,4]. Our numerical experiments covering around 30 000 synthesis instances per each of the three universal quantum bases considered have not produced a single instance that would violate an underlying conjecture.

It is important to point out the general and largely open problem of synchronizing probabilistic circuits with state distillation protocols. We have estimated the costs of the logical PQF circuits in terms of logical basis gates. In many architectures, however, the key basis gates, such as $\pi/8$, $\pi/12$ gates or the V gate, require distilled magic states that are commonly prepared external to the primary circuit, in either an offline or online fashion. The demand of PQF circuits on distilled states is probabilistic and, although the probability of excessive demand is very low, in the worst-case scenario

this may add latency to the execution of quantum algorithms. A direction for future work is to determine optimization and compilation strategies that minimize the resource requirements of PQF circuits in the context of a logical quantum circuit. Such strategies could include considering the circuit depth and width costs of PQF solutions when layered with quantum error correction as well as state distillation protocols.

Recent methods have addressed the distillation of non-Clifford states and Fourier states, and provide a starting point for research [27,28] on other possible universal bases and their fault-tolerant constructions. Consideration of generalization to qudit computation models is also an avenue for future exploration. Finally, formal, rigorous proofs of the underlying conjectures is another important direction.

APPENDIX A: INFORMATION-THEORETIC BOUNDS

The density with which cyclotomic rationals are distributed imposes information-theoretic limits on how much we can reduce the expected T count of our nondeterministic solutions compared to the T count of deterministic, unitary solutions. Note that the analysis applies equally well to both PQF and RUS methods.

Let us assume, temporarily, that for $z \in \mathbb{Z}[\omega]$ and $L = \lceil \log_2(|z|^2) \rceil$, the norm equation $|y|^2 = 2^L - |z|^2$ is solvable. By definition of L , $|z|^2 \leq 2^L$. We know that the optimal T count of a single-qubit unitary circuit implementing a matrix of the form of Eq. (3) is $t = 2L$ or $t = 2L - 2$.

In either case we note that $|z|^2 = O(2^{t/2})$ and $|z|^4 = O(2^t)$. We also note that given an upper bound b on the absolute value of cyclotomic integer, there are no more than $O(b^4)$ cyclotomic integers under this bound. Thus, we conclude that there are no more than $O(2^t)$ cyclotomic integers z for which the matrix of the form of Eq. (3) may exist and be implemented at T count t or less.

It follows that there are at most $O(2^t)$ unimodular cyclotomic rationals on the unit circumference for which our RUS circuit can be built with design cost of T count = t or less. Therefore, there exists a constant K such that for $\epsilon < K \times 2^{-t}$ there is an arc of the unit circumference of length 2ϵ that does not contain any such cyclotomic rational. If θ_* is the angle in the center of such an arc, then the rotation $R_z(\theta_*)$ cannot be implemented by any of our circuits with design cost of T count = t or less.

Conversely, $\epsilon \geq K \times 2^{-t}$ is the necessary condition for any axial rotation to be implementable by one of our circuits with design cost of T count = t or less. This necessary condition is equivalent to

$$t \geq \log_2(1/\epsilon) + \log_2(K), \quad (\text{A1})$$

which is a specific lower bound on the design cost given by the T count of our solution.

The derivation of the above lower bound is specific to our PQF and RUS designs. As follows from Corollary 14, our existing PQF protocol for single-qubit decomposition based on the PSLQ integer relation algorithm is within $O\{\ln[\ln(1/\epsilon)]\}$ from this bound. Deriving a uniform lower bound under more general assumptions would be a worthwhile problem for future research.

APPENDIX B: DETAILS ON THE NORM EQUATION IN $\mathbb{Z}[\omega]$

This section combines the claims for $\omega = e^{2\pi i/m}$ for $m = 8$ and 12 . We reintroduce $\rho = \omega + \omega^*$. Recall that the real-valued ring $\mathbb{Z}[\rho]$ is a unique factorization ring. That is, any of its elements can be factored into a product of prime algebraic integers and at most one unit. The primary category of right-hand-side values for which Eq. (6) is easily solvable would then be the set of algebraic integer primes.

The equation is easily solvable for the following kinds of prime right-hand sides (cf. [19]):

(1) $\xi = a + b\rho$, $\xi > 0$ and $p = \xi\xi^*$ is a positive rational prime number with $p \equiv 1 \pmod{m}$;

(2) ξ is a rational prime number and $\xi \not\equiv -1 \pmod{m}$.

We call an algebraic integer prime belonging to one of these two classes a “good” prime.

For a composite ξ we consider a *limited* factorization of the right-hand side to preserve efficiency. To this end, we precompute a set $S_{\text{prime}} \subset \mathbb{Z}[\rho]$ of small prime elements and consider factorizations of the form $\xi = \xi_1^{a_1} \dots \xi_r^{a_r} \eta$, where $\xi_1, \dots, \xi_r \in S_{\text{prime}}$ and η passes a primality test. Equation (6) is efficiently solvable if η is a good prime and for $i = 1, \dots, r$, ξ_i is a good prime or a_i is even.

Example 15. For $m = 8$, $|y|^2 = \xi = 1\,270\,080 + 211\,680\sqrt{2}$ is efficiently solvable since $\xi = 2^5 3^3 5^7 (2 + \sqrt{2})(5 - 2\sqrt{2})$.

Note $p = (5 - 2\sqrt{2})(5 + 2\sqrt{2})^* = 17 \equiv 1 \pmod{8}$. The only “bad” prime in the above factorization is 7 but it appears as an even power. We remark that the cyclotomic integer z coming from the cyclotomic rational approximation of $e^{i\theta}$ is not unique. In fact, it is defined up to an arbitrary real-valued factor $r \in \mathbb{Z}[\rho]$. For any such r , $(rz)^*/(rz)$ is identical to z^*/z . However, the norm equation $|y|^2 = 2^L - |rz|^2$ can and will change quite dramatically.

When drawing r randomly from a subset of $\mathbb{Z}[\rho]$ one might try and estimate the chance that the equation $|y|^2 = 2^L - |rz|^2$ turns out to be solvable for a random r . This is an example of an open and likely very hard number-theory problem. We will not attempt to solve it here and will instead rely on a conjecture that the “lucky” values of r are reasonably dense in $\mathbb{Z}[\rho]$.

APPENDIX C: EXACTLY REPRESENTABLE SINGLE-QUBIT CIRCUITS IN CLIFFORD + $\pi/12$

We use the notation $\omega = \omega_{12} = e^{i\pi/6}$ in this section. In this and subsequent sections, we also use a shorthand notation for single-qubit controlled phase gate. Given $\phi \in \mathbb{C}, |\phi| = 1$ is a phase factor, the *controlled phase gate* $\Lambda(\phi)$ is simply $\begin{bmatrix} 1 & 0 \\ 0 & \phi \end{bmatrix}$. In particular, the $\pi/12$ gate $K = \Lambda(\omega_{12})$.

The single-qubit Clifford + $\pi/12$ group is generated by the Hadamard gate H and the $\pi/12$ gate K . We note that $\omega^3 = i$ and therefore the common phase gate $S = K^3$ is in the circuit group, as is, by closure, the entire single-qubit Clifford group.

Any Clifford + $\pi/12$ circuit can be expressed as a product of syllables of the form $K^k H$, where $|k| < 6$, up to a possible global phase factor. A slightly deeper analysis reveals that

we can rewrite a circuit to enforce $k = \pm 1, \pm 2$ in all interior syllables, but this is not very important in this section. The important part is that the $K^k H$ syllable is a Clifford gate for $k \in \{0, \pm 3, \pm 6, \pm 9\}$ and has zero K count.

We assume that the implementation cost of gates of the form $K^k, k \notin \{0, \pm 3, \pm 6, \pm 9\}$ is the same and that it is significantly higher than the cost of a Clifford gate. This implies that the K cost of a circuit composed of $K^k H$ syllables is upper bounded by the number of syllables with $k \notin \{0, \pm 3, \pm 6, \pm 9\}$.

Consider the ring of cyclotomic integers $\mathbb{Z}[\omega]$. Any ω^k is a cyclotomic integer and $H = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & -1 \\ 1 & -1 \end{bmatrix}$, where ± 1 are in $\mathbb{Z}[\omega]$ and $\frac{1}{\sqrt{2}}$ is not in $\mathbb{Z}[\omega]$. Clearly, a finite product of the $K^k H$ syllables evaluates to a unitary matrix of the form

$$\frac{1}{\sqrt{2}^L} \begin{bmatrix} z & -y^* \omega^\ell \\ y & z^* \omega^\ell \end{bmatrix}, \quad (\text{C1})$$

$y, z \in \mathbb{Z}[\omega], \ell, L \in \mathbb{Z}$.

Lemma 16. A Clifford + $\pi/12$ circuit that evaluates to a unitary in the form of Eq. (C1) with $L = 0$ has K cost 0 or 1.

Proof. The unitary condition of the matrix in Eq. (C1) with $L = 0$ means $|z|^2 + |y|^2 = 1$. Since y, z are algebraic integers, either $|z| = 1, |y| = 0$, or $|z| = 0, |y| = 1$. By standard algebraic units argument, if $x \in \mathbb{Z}[\omega]$ and $|x| = 1$ then, $x = \omega^k, k \in \mathbb{Z}$. In the case $|z| = 1$ and $z = \omega^k$, the unitary in the form of Eq. (C1) is $\omega^k \Lambda(\omega^{\ell-2k})$. As per the assumptions we have adopted above, the $\pi/12$ cost of the latter is either 0 or 1. The case of $|z| = 0$ is reduced to the case of $|z| = 1$ by premultiplying the subject unitary times $X = HZH = H\Lambda(\omega^6)H$. By our convention, the latter has zero K cost and does not affect the K cost of the resulting circuit. ■

Lemma 17. Consider a unitary 2-vector of the form $v = \frac{1}{\sqrt{2}^L} (z, y)^T, y, z \in \mathbb{Z}[\omega], L \in \mathbb{Z} : |y|^2 + |z|^2 = 2^L$. A Clifford + $\pi/12$ circuit c with K cost at most $L + 1$ can be algorithmically found such that $c v = (1, 0)^T$.

Proof. This rather technical lemma is inspired by the “column lemma” from [26]. The proof is by induction in L . The base of the induction is $L = 0$, and the claim has been already established in the proof of Lemma 16.

Consider the subject vector with $L > 0$. The main step is to algorithmically find a short circuit c with K cost at most 1 such that $v' = c v$ is or the form $v' = \frac{1}{\sqrt{2}^{L'}} (z', y')^T$, where $L' < L$. Then, the desired circuit will be generally of the form $H\Lambda(\omega^k)$, except for one special case where it will be a global phase. We generally attempt to find $k \in \mathbb{Z}$ such that all the integer coefficients of the algebraic integers $z \pm \omega^k y$ are even. If we have succeeded in finding such a k then

$$\begin{aligned} H\Lambda(\omega^k)v &= \frac{1}{\sqrt{2}^{L+1}} (z + \omega^k y, z - \omega^k y)^T \\ &= \frac{1}{\sqrt{2}^{L-1}} [(z + \omega^k y)/2, (z - \omega^k y)/2]^T \end{aligned}$$

and we have succeeded in reducing the denominator exponent.

In order to develop a method for finding the desired k , consider the parity morphism

$$\begin{aligned} \mu &: \mathbb{Z}[\omega] \rightarrow \mathbb{Z}_2[\omega], \\ \mu &: a\omega^3 + b\omega^2 + c\omega + d \\ &\mapsto (a \bmod 2)\omega^3 + (b \bmod 2)\omega^2 + (c \bmod 2)\omega + (d \bmod 2). \end{aligned}$$

All coefficients of $z \pm \omega^k y$ are even if and only if $0 = \mu(z \pm \omega^k y) = \mu(z) \oplus \omega^k \mu(y)$ if and only if $\mu(z) = \omega^k \mu(y)$, which is going to be the desired property below.

Consider the action of the 12-element group $\{\omega^k\}$ on $\mathbb{Z}_2[\omega]$ by multiplication. Since $\omega^6 = -1$ and $-1 = 1 \bmod 2$ the subgroup $\{1, -1\}$ acts trivially on $\mathbb{Z}_2[\omega]$ and the action of the 6-element factor group $\{\omega^k\}/\{1, -1\}$ is well defined.

By direct computation we established that the 16-element set $\mathbb{Z}_2[\omega]$ is partitioned into four orbits of this action. The orbit of zero O_0 consists of just zero. The orbit of $1 + \omega^3 = 1 + i$, O_3 consists of three elements and the orbits O_1 and O_2 of 1 and $1 + \omega$, respectively, consist of six elements each.

It is important to note that the function $N_2 : x \mapsto \mu(|x|^2)$ is constant on each of the orbits. More specifically, $N_2(O_0) = N_2(O_3) = 0$, $N_2(O_1) = 1$, $N_2(O_2) = \omega^3$. The final key remark is that the unitary constraint of the vector v implies $N_2(z) \oplus N_2(y) = N_2(2^L) = 0$ and therefore $N_2(z) = N_2(y)$.

We proceed by case distinction.

(1) (0,0) Case of O_0 : If both $\mu(z)$ and $\mu(y)$ belong to O_0 then all the integer coefficients of y and z are already even and we do not need to do any transformations in order to reduce the vector.

(2) (1,2) Cases of O_1 and O_2 : If $\mu(z)$ belongs to either of the two orbits, then $\mu(y)$ must belong to the same orbit [since we have established $N_2(z) = N_2(y)$]. Therefore, there exists k such that $\mu(z) = \omega^k \mu(y)$ which is what we were looking for.

(3) (3,3) Case of O_3 : If both $\mu(z)$ and $\mu(y)$ belong to O_3 , then again there exists k such that $\mu(z) = \omega^k \mu(y)$.

(4) (3,0): This is the only remaining case.

If one and only one of the $\mu(z)$, $\mu(y)$ belongs to O_3 , then the other one must belong to O_0 [since these are the only two orbits with $N_2(\text{orbit}) = 0$]. Assume, w.l.o.g. that $\mu(y) = 0$.

This case needs to be treated differently from the general context. First, we note that, since $\mu(z) \in O_3$ there exists a k such that $\mu(\omega^k z) = 1 + \omega^3 = 1 + i$. Next, we note that the global phase operator $(1 + i)/\sqrt{2}I_2$ is in the Clifford group and that $\mu[(1 + i)^2] = \mu(2\omega^3) = 0$. Therefore, by multiplying the vector v times the global phase $\omega^k(1 + i)/\sqrt{2}I_2$ we obtain a vector, where all the integer coefficients of both components are even. We then reduce this latter vector to one of the form $1/\sqrt{2}^{L-1} w$.

This case concludes the induction step. \blacksquare

Corollary 18. Unitary of the form in Eq. (C1), where $y, z \in \mathbb{Z}[\omega]$, $L, k \in \mathbb{Z}$, can be represented exactly and algorithmically by a Clifford + $\pi/12$ circuit of $\pi/12$ count at most $L + 2$.

Proof. Consider a Clifford + $\pi/12$ circuit c of $\pi/12$ count at most $L + 1$ that reduces the first column of the matrix in Eq. (C1) to $(1, 0)^T$. Consider the unitary value of $c^\dagger = c^\dagger I_2$. Since c^\dagger maps $(1, 0)^T$ into the first column of Eq. (C1) it maps $(0, 1)^T$ into a unitary vector that is Hilbert orthogonal to that first column. Thus, $c^\dagger(0, 1)^T$ is proportional

to $\frac{1}{\sqrt{2}^L}(-y^* \omega^\ell, z^* \omega^\ell)^T$ with a unit coefficient from $\mathbb{Z}[\omega]$. Therefore, we can algorithmically find an integer k such that the unitary in Eq. (C1) is exactly equal to the value of $c^\dagger \Lambda(\omega^k)$. Since the $\pi/12$ count of the $\Lambda(\omega^k)$ is at most 1 the corollary follows. \blacksquare

APPENDIX D: APPROXIMATION OF REAL NUMBERS BY NUMBERS FROM $\mathbb{Z}[\sqrt{3}]$

This section is a direct extension of Sec. 5 in [3] to the $\mathbb{Z}[\sqrt{3}]$ ring. Recall that the fundamental Galois automorphism of that ring extends $\bullet : \sqrt{3} \mapsto (-\sqrt{3})$. The following is an analog of Lemma 17 from [3]:

Lemma 19. Let $[x_0, x_1]$ and $[y_0, y_1]$ be closed intervals of real numbers. Let $\delta = x_1 - x_0$ and $\Delta = y_0 - y_1$, and assume $\delta \Delta \geq (2 + \sqrt{3})^2$. Then, there exists at least one $\alpha = a + b\sqrt{3} \in \mathbb{Z}[\sqrt{3}]$ such that $\alpha \in [x_0, x_1]$ and $\alpha^\bullet = a - b\sqrt{3} \in [y_0, y_1]$. Moreover, there is an efficient algorithm for computing such a and b .

The proof is almost identical to the proof of the lemma for $\mathbb{Z}[\sqrt{2}]$ with the obvious replacements of $\sqrt{2}$ by $\sqrt{3}$ and of the unit $\lambda = 1 + \sqrt{2}$ by the unit $\nu = 2 + \sqrt{3}$.

APPENDIX E: APPROXIMATING SINGLE-QUBIT CIRCUITS IN CLIFFORD + $\pi/12$

This section is a direct extension of Sec. 6 in [3] to the $\mathbb{Z}[\omega = e^{i\pi/6}]$ ring. We prove that an axial rotation $\Lambda(e^{i\theta})$ can be algorithmically approximated to any desired precision $\varepsilon > 0$ by a Clifford + $\pi/12$ circuit with $\pi/12$ count of at most $2 \log_{\sqrt{2}}(1/\varepsilon) + C$, where $C = 3/2 + \log_{\sqrt{2}}(2 + \sqrt{3})$. Recall that $\sqrt{3} = 2\omega - \omega^3$ and $i = \omega^3$ and consider the subring $\mathbb{Z}[\sqrt{3}][i] \subset \mathbb{Z}[\omega]$. Let $\theta \in \mathbb{R}$ and $\varepsilon > 0$ be fixed.

Definition 20. Consider some $u = (a + b\sqrt{3}) + (c + d\sqrt{3})i \in \mathbb{Z}[\sqrt{3}][i]$. Complex number $u/\sqrt{2}^k$, $k \in \mathbb{Z}$ is called a feasible candidate at round k for (θ, ε) if

$$(1) |u^\bullet|^2 \leq 2^k;$$

$$(2) |u|^2 \leq 2^k \text{ and } \text{Re}(ue^{i\theta/2}) \geq (1 - \varepsilon^2)\sqrt{2}^k.$$

Theorem 21. Let $\varepsilon > 0$ and $\theta \in \mathbb{R}$ be fixed and let $k \geq C + \log_{\sqrt{2}}(1/\varepsilon)$, where $C = 1/2 + \log_{\sqrt{2}}(2 + \sqrt{3})$. Then, there exists a set of at least $n = \lfloor 2\sqrt{2}/\varepsilon \rfloor$ feasible candidates at round k for (θ, ε) . Moreover, there is an efficient algorithm for generating a sequence of random candidates from this set.

Proof. First note that $k \geq C + \log_{\sqrt{2}}(1/\varepsilon)$ implies $2^k \geq \sqrt{2}(2 + \sqrt{3})^2/\varepsilon^2$. Define $\delta = \sqrt{2}^k \varepsilon^2$ and $\Delta = \sqrt{2}^{k+1}$. and observe that $\delta \Delta \geq (2 + \sqrt{3})^2$ so that the criterion of Lemma 19 is satisfied for (δ, Δ) . For convenience we assume w.l.o.g. that $-\pi/2 \leq \theta \leq \pi/2$.

Using the same geometric argument as in proof of Theorem 22 in [3] we observe that condition 2 of Definition 20 defines a meniscus shape R_ε on the complex plane. If we parametrize the plane with $x + yi$, $x, y \in \mathbb{R}$ we observe that there is a vertical segment $[y_{\min}, y_{\max}]$ such that $y_{\max} - y_{\min} \geq \sqrt{2}\varepsilon$ and such that for any $y' \in [y_{\min}, y_{\max}]$ the intersection of the horizontal line $\{x + y'i\}$ with the meniscus R_ε is a segment of length at least $\varepsilon^2/2$.

Introduce $n = \lfloor 2\sqrt{2}/\varepsilon \rfloor$. We now partition the segment $[y_{\min}, y_{\max}]$ at points $y_j = j/n(y_{\max} - y_{\min}) + y_{\min}$, $j = 0, \dots, n$. By design, $y_{j+1} - y_j > \varepsilon^2/2$. Consider closed subintervals $I_j = [y_j, y_j + \varepsilon^2/2]$, $j = 0, \dots, n-1$, that are nonoverlapping subintervals of the $[y_{\min}, y_{\max}]$. First, we find $\beta_j \in \mathbb{Z}[\sqrt{3}]$ such that $\beta_j \in [\sqrt{2}^k y_j, \sqrt{2}^k (y_j + \varepsilon^2/2)]$ and $\beta_j^* \in [-\sqrt{2}^{k-1}, \sqrt{2}^{k-1}]$. This can be done algorithmically because $|\lfloor -\sqrt{2}^{k-1}, \sqrt{2}^{k-1} \rfloor| \sqrt{2}^k \varepsilon^2/2 \geq (2 + \sqrt{3})^2$.

Let $H_j = R_\varepsilon \cap \{y = \beta_j/\sqrt{2}^k\}$. As we have noted, the length of H_j is at least $\varepsilon^2/2$. Now find $\alpha_j \in \mathbb{Z}[\sqrt{3}]$ such that $\alpha_j \in \sqrt{2}^k H_j$ and $\alpha_j^* \in [-\sqrt{2}^{k-1}, \sqrt{2}^{k-1}]$. This can be done algorithmically for the same reason as above.

Clearly, $(\alpha_j + \beta_j i)/\sqrt{2}^k$ is a feasible candidate at round k for (θ, ε) and it is distinct from any other such candidate $(\alpha_{j'} + \beta_{j'} i)/\sqrt{2}^k$, $j \neq j'$. By randomly selecting an integer $0 \leq j < n$ without replacement, we now can algorithmically generate a unique random feasible candidate as claimed. ■

We now discuss a conjecture regarding solvability of a norm equation that is needed for expanding a feasible candidate $z/\sqrt{2}^L, z \in \mathbb{Z}[\omega]$ into a unitary matrix of the form in Eq. (C1). Such expansion exists if and only if the norm equation $|y|^2 = \xi = 2^L - |z|^2$ can be solved for $y \in \mathbb{Z}[\omega]$. In Appendix B, we have defined the notion of easily solvable norm equation and also built up sufficient intuition the the effect that easily solvable norm equations are not uncommon. They are more common than the prime numbers with the additional property $p \equiv 1 \pmod{12}$ among the integers. It is well known that in a segment of the form $[B/2, B]$, where B is sufficiently large, the density of such prime numbers is in $\Omega[1/\ln(B)]$.

Further steps in the single-qubit circuit synthesis rely on the following conjecture (of the type that is now becoming common in circuit synthesis):

Conjecture 22. For small enough values of $\varepsilon > 0$ and L in $\Omega[\ln(1/\varepsilon)]$ it suffices to inspect $O[\ln(2^L)] = O(L)$ feasible candidates $z/\sqrt{2}^L$ for (θ, ε) in order to find at least one such candidate for which the norm equation $|y|^2 = 2^L - |z|^2$ is easily solvable over $\mathbb{Z}[\omega]$.

Rigorous proof of this conjecture may be a hard number-theory problem. At this time, however, we have ample numeric evidence for the conjecture for a range of ε down to 10^{-100} .

Assuming this conjecture we can claim the following:

Theorem 23. Let θ be a fixed angle. There exists a synthesis algorithm with probabilistically polynomial classical runtime that solves the following problem: For a small enough value of $\varepsilon > 0$ find a unitary ancilla-free Clifford + $\pi/12$ circuit with $\pi/12$ count smaller than $2 \log_2(1/\varepsilon) + K$ {where $K = \lceil 5/2 + 2 \log_2(2 + \sqrt{3}) \rceil$ } that represents the axial rotation $\Lambda(e^{i\theta})$ to absolute precision ε .

Proof. Given (θ, ε) , Theorem 21 algorithmically defines a set of feasible candidates $z/\sqrt{2}^L$ or cardinality $\Omega(1/\varepsilon)$. For all these candidates $L \leq 2 \log_2(1/\varepsilon) + K - 2$ and is in $O[\ln(1/\varepsilon)]$. The outer loop of the desired algorithm randomly samples feasible candidates from the above set without replacement. By Conjecture 22, with arbitrarily high probability the algorithm finds a feasible candidate with an easily solvable norm equation after $O(L)$ trials.

Let $z/\sqrt{2}^L$ be such candidate and $y \in \mathbb{Z}[\omega]$ be a solution of the norm equation $|y|^2 = 2^L - |z|^2$. Then, the unitary

$$\frac{1}{\sqrt{2}^L} \begin{bmatrix} z & -y^* \\ y & z^* \end{bmatrix} \quad (\text{E1})$$

is an ε approximation of the rotation $\Lambda(e^{i\theta})$. Per Corollary 18, this unitary can be exactly represented by a Clifford + $\pi/12$ circuit with $\pi/12$ count at most $L + 2$, and the theorem follows. ■

APPENDIX F: RUNTIME PERFORMANCE EVALUATION

In this section, we report empirical results for upper bounds on the compilation time of the PQF synthesis algorithm, focusing primarily on the more expensive stages 1 and 2, for each universal gate set presented.

In the case of synthesis over V basis, phase-factor approximation is performed using a simple continued fraction algorithm and has trivial cost compared to the stage 2 cost. In the cases of Clifford + T and Clifford + $\pi/12$ bases, we have implemented the PSLQ algorithm [22] in *Mathematica*.

The main theorem in Ref. [20] states that if exact integer relations between the subject real values exist and M is the minimum norm of such an integer relation, then the PSLQ algorithm terminates after a number of iterations bounded by $O[\ln(M)]$. Both the theorem and the proof can be customized to find approximate integer relations; the modification then states that if M_ε is the minimum size of an integer vector a such that $|a \cdot x| < \varepsilon$, then the modified algorithm terminates after a number of iterations bounded by $O[\ln(M_\varepsilon)]$. Since in the case of both T and $\pi/12$ bases, $M_\varepsilon = O(\varepsilon^{-1/4})$, the bound on the number of iterations before termination is linear in $\ln(1/\varepsilon)$.

The Bertok implementation of the PSLQ algorithm [22] appears to be asymptotically optimal. In our experiments on 1000 random target angles, the number of PSLQ iterations scales on average as $1.16 \log_2(1/\varepsilon)$ for the Clifford + T basis and as $1.04 \log_2(1/\varepsilon)$ for the Clifford + $\pi/12$ basis. The standard deviation on the number of iterations scales roughly as $\log_2[\log_2(1/\varepsilon)]$.

The practical cost of stage 1 becomes quadratic in $\ln(1/\varepsilon)$ for $\varepsilon < 10^{-15}$ when measured in common arithmetic operations native to a classical computer because the PSLQ algorithm requires variable precision floating point arithmetic with precision tightening as $O(\varepsilon)$. Once the required precision exceeds the available machine precision, software simulation of variable mantissa arithmetic becomes necessary, causing a one-time drop in speed and a subsequent quadratic trend in compilation cost.

The compilation time for each stage 2 round is roughly proportional to the number of candidate modification factors evaluated until an easy solution to a suitable norm equation is found. Each candidate factor is generated by an appropriate enumerator, then the corresponding norm equation is tested for easy solvability. The cost of generating a candidate is trivial compared to the cost of analyzing the norm equation. The latter cost is in principle similar to the cost of testing an integer for smoothness; however, in our prototype implementation we simply relied on the *Mathematica* FACTOR INTEGER function,

time constrained to one quarter of a second. The cost of analyzing the factors is trivial compared to the cost of the factorization.

We can upper bound the compilation time of stage 2 by $t/4$ seconds, where t is the number of the candidate factors needed for termination. While t should be expected to scale like $O[\ln(1/\varepsilon)]$ with precision, we find in practice that t and hence the time required for stage 2 is a strongly stochastic variable. While the time is practical in all cases, outlying cases might require an order of magnitude more candidates to terminate than typical cases.

For the V basis, the mean expected value of t scales as $1.2 + 0.36 \log_5(1/\varepsilon)$. The maximum number of candidates scales as $8.8 + 4 \log_5(1/\varepsilon)$. Somewhat surprisingly, the value t required for stage 2 in the case of Clifford + T and Clifford + $\pi/12$ bases shows very little correlation with the target precision (insignificant correlation coefficient) with a mean expectation of t around 2.2 for the Clifford + T basis and around 2.1 for the Clifford + $\pi/12$ basis. The expected maximum t also appears

uncorrelated and is around 23 for the Clifford + T basis and 22 for the Clifford + $\pi/12$ basis.

Thus, while t has single digits for the majority of (θ, ε) pairs it occasionally grows significantly in outlying cases. The outlying cases, however, still remain practically manageable and finish in seconds due to throttling of integer factorization. A conceptual explanation of the apparent stochastic behavior remains to be found. It may be related to the apparent fractal structure of the set of solvable norm equations.

To summarize, we find that the expected compilation time for the PQF algorithm for all three bases is linear in $\ln(1/\varepsilon)$ at coarse precisions and becomes quadratic in $\ln(1/\varepsilon)$ at finer precisions. The compilation time can occasionally spike for outlying (θ, ε) pairs due to fluctuations in the required number of modifier candidates; this translates into only seconds when run on a common desktop computer. We believe the performance speed can be further optimized by reimplementing the algorithm in a fully compiled language.

-
- [1] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, UK, 2000).
- [2] C. Dawson and M. Nielsen, The Solovay-Kitaev algorithm, *Quantum Inf. Comput.* **6**, 81 (2006).
- [3] P. Selinger, Efficient Clifford + T approximation of single-qubit operators, [arXiv:1212.6253](https://arxiv.org/abs/1212.6253).
- [4] N. Ross and P. Selinger, Optimal ancilla-free Clifford + T approximation of z -rotations, [arXiv:1403.2975](https://arxiv.org/abs/1403.2975).
- [5] V. Kliuchnikov, D. Maslov, and M. Mosca, Practical approximation of single-qubit unitaries by single-qubit quantum Clifford and T circuits, [arXiv:1212.6964](https://arxiv.org/abs/1212.6964) [*IEEE Trans. Comput.* (2015), doi: 10.1109/TC.2015.2409842].
- [6] V. Kliuchnikov, Synthesis of unitaries with Clifford + T circuits, [arXiv:1306.3200](https://arxiv.org/abs/1306.3200).
- [7] A. Bocharov, M. Roetteler, and K. M. Svore, Efficient synthesis of universal Repeat-Until-Success circuits, *Phys. Rev. Lett.* **114**, 080502 (2015).
- [8] A. Paetznick and K. M. Svore, Repeat-Until-Success: Non-deterministic decomposition of single-qubit unitaries, *Quantum Inf. Comput.* **14**, 1277 (2014).
- [9] A. Bocharov, Y. Gurevich, and K. M. Svore, *Phys. Rev. A* **88**, 012313 (2013).
- [10] A. Blass, A. Bocharov, and Y. Gurevich, Optimal Ancilla-free Pauli+ V Circuits for Axial Rotations, [arXiv:1412.1033](https://arxiv.org/abs/1412.1033) (2014).
- [11] N. J. Ross, Optimal ancilla-free Pauli + V approximation of z -rotations, [arXiv:1409.4355](https://arxiv.org/abs/1409.4355).
- [12] S. X. Cui and Z. Wang, Universal quantum computation with metaplectic anyons, [arXiv:1405.7778](https://arxiv.org/abs/1405.7778).
- [13] V. Kliuchnikov, D. Maslov, and M. Mosca, Fast and efficient exact synthesis of single qubit unitaries generated by Clifford and T gates, *Quantum Inf. Comput.* **13**, 607 (2013).
- [14] J. Norris, *Markov Chains* (Cambridge University Press, Cambridge, UK, 1998).
- [15] A. Bocharov and K. M. Svore, A depth-optimal canonical form for single-qubit quantum circuits, *Phys. Rev. Lett.* **109**, 190501 (2012).
- [16] K. Matsumoto and K. Amano, Representation of quantum circuits with clifford and $\pi/8$ gates, [arXiv:0806.3834](https://arxiv.org/abs/0806.3834) (2008).
- [17] B. Giles and P. Selinger, Remarks on Matsumoto and Amano's normal form for single-qubit Clifford + T operators, [arXiv:1312.6584](https://arxiv.org/abs/1312.6584).
- [18] M. B. Hastings and A. Geller, Reduced space-time and time costs using dislocation codes and arbitrary ancillas, [arXiv:1408.3379](https://arxiv.org/abs/1408.3379).
- [19] L. Washington, *Introduction to Cyclotomic Fields* (Springer, New York, 1997).
- [20] H. Ferguson and D. Bailey, A polynomial time, numerically stable integer relation algorithm, <http://crd-legacy.lbl.gov/~dhbailey/dhbpapers/pslq.pdf>.
- [21] W. Schmidt, *Diophantine Approximation* (Springer, Berlin, 1980).
- [22] P. Bertok, PSLQ integer relation algorithm implementation, <http://library.wolfram.com/infocenter/MathSource/4263>.
- [23] S. Bravyi and A. Y. Kitaev, Universal quantum computation with ideal clifford gates and noisy ancillas, *Phys. Rev. A* **71**, 022316 (2005).
- [24] A. Meier, B. Eastin, and E. Knill, Magic-state distillation with the four-qubit code, [arXiv:1204.4221](https://arxiv.org/abs/1204.4221).
- [25] S. Bravyi and J. Haah, Magic state distillation with low overhead, *Phys. Rev. A* **86**, 052329 (2012).
- [26] B. Giles and P. Selinger, Exact synthesis of multiqubit Clifford+ T circuits, *Phys. Rev. A* **87**, 032332 (2013).
- [27] G. Duclos-Cianci and D. Poulin, Reducing the quantum-computing overhead with complex gate distillation, *Phys. Rev. A* **91**, 042315 (2015).
- [28] G. Duclos-Cianci and K. M. Svore, A state distillation protocol to implement arbitrary single-qubit rotations, *Phys. Rev. A* **88**, 042325 (2013).