

Simulating Quantum Computation: How Many “Bits” for “It”?


Michael Zurel^{1,2,*}, Cihan Okay³, and Robert Raussendorf^{2,4}

¹Department of Physics & Astronomy, *University of British Columbia, Vancouver, British Columbia V6T 1Z1, Canada*

²Stewart Blusson Quantum Matter Institute, *University of British Columbia, Vancouver, British Columbia V6T 1Z4, Canada*

³Department of Mathematics, *Bilkent University, 06800 Ankara, Turkey*

⁴Institute of Theoretical Physics, *Leibniz Universität Hannover, 30167 Hannover, Germany*

 (Received 28 January 2024; revised 22 April 2024; accepted 25 July 2024; published 3 September 2024)

A recently introduced classical simulation method for universal quantum computation with magic states operates by repeated sampling from probability functions [M. Zurel *et al.* PRL 260404 (2020)]. This method is closely related to sampling algorithms based on Wigner functions, with the important distinction that Wigner functions can take negative values obstructing the sampling. Indeed, negativity in Wigner functions has been identified as a precondition for a quantum speed-up. However, in the present method of classical simulation, negativity of quasiprobability functions never arises. This model remains probabilistic for all quantum computations. In this paper, we analyze the amount of classical data that the simulation procedure must track. We find that this amount is small. Specifically, for any number n of magic states, the number of bits that describe the quantum system at any given time is $2n^2 + O(n)$.

DOI: [10.1103/PRXQuantum.5.030343](https://doi.org/10.1103/PRXQuantum.5.030343)

I. INTRODUCTION

In an article entitled “Information, Physics, Quantum: The Search for Links” [1], J.A. Wheeler argued that quantum physics required a new perspective on reality based on information theoretic concepts. He wrote: “No element in the description of physics shows itself as closer to primordial than the elementary quantum phenomenon, that is, the elementary device-intermediated act of posing a yes-no physical question and eliciting an answer or, in brief, the elementary act of observer-participancy. Otherwise stated, every physical quantity, every it, derives its ultimate significance from bits, binary yes-or-no indications, a conclusion we epitomize in the phrase, *it from bit*.”

A prototypical realization of this view on physics has been provided in the description of quantum computation with magic states (QCM) through the Λ polytopes [2], in which the quantum phenomena at hand are reproduced, without any deviation or approximation, by repeated sampling processes dependent on and producing bit strings. The “*it*” in this case is universal quantum computation; hence all nonrelativistic quantum mechanics in finite-dimensional Hilbert spaces. The “*bits*” represent the binary

outcomes of Pauli measurements and the labels of the vertices of the Λ polytopes in the repeated sampling. There are finitely many such vertices for any number n of magic states.

This description of quantum computation with magic states grew out of the analysis of Wigner function negativity as a precondition for a quantum computational speed-up, a research program that started with Refs. [3,4]. Specifically, in Ref. [4], it was established that negativity in Gross’s Wigner function [5,6] is required for a quantum speed-up, under the condition that the Hilbert-space dimension is odd. Analogous results were subsequently established in even dimension, specifically for rebits [7] and qubits [8]. However, in the end it has turned out that once sufficiently general (quasi)probability functions are admitted, there is no need for any negativity at all [2,9]. Universal quantum computation in the magic state model can be described by repeated sampling from a generalized phase space the points of which are labeled by the vertices of the Λ polytopes. This process essentially resembles a random walk, with the complication that the transition function changes from one time step to the next and can depend on the prior sampling history. A summary of this sampling process is given in Sec. III; see in particular Theorem I and Algorithm 1.

The question of interest for the present work is how much classical information the simulation of quantum computations must track, i.e., *What is the length of the bit strings produced in the simulation?* For example, it turned out that those bit strings were very long—say,

*Contact author: mzurel@phas.ubc.ca

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

exponentially long in the number n of magic states—this would provide a convenient explanation for the hardness of classical simulation of universal quantum computation using Λ polytopes. If the information storage itself is inefficient, so is the processing. However, this is not what we find. We find that the bit strings are short. Specifically, they are of length $O(n^2)$. Thus, simulation of universal quantum computation based on Λ polytopes is a *small-data problem*. The presumed hardness of this simulation must come from the computational hardness of the sampling processes involved, not from moving around large amounts of data [10].

Naively, the lengths of the bit strings labeling the vertices of Λ_n are upper bounded by $\log_2(|\mathcal{V}_n|)$, where $|\mathcal{V}_n|$ is the size of the generalized phase space, i.e., the number of vertices of Λ_n . To date, an upper bound and a lower bound are known for this quantity, namely,

$$\frac{n^2}{2} \leq \log_2(|\mathcal{V}_n|) \leq 4^n n^2.$$

The lower bound is by Karanjai *et al.* [11]. The upper bound comes from the upper-bound theorem of polytope theory [12] (for details, see Appendix A).

The gap between the bounds is extremely wide, consistent with both efficient and inefficient storage of the bit strings. Numerical studies suggest that the number of phase-space points is growing very rapidly with n ; for $n = 1$, the number of phase-space points is 8, for $n = 2$, it is 22 320, and for $n > 2$, we do not know the precise numbers but the estimate for $n = 3$ is already huge.

The following simple insight is crucial for establishing our main result: in the QCM model, for any fixed value n , all quantum computations start in the *same* magic state $|M\rangle^{\otimes n}$. Therefore, the question of interest for classical simulation of QCM using Λ polytopes is not “What is the size of the phase space \mathcal{V}_n ?” but, rather, “What is the size of the region of \mathcal{V}_n that can be reached from the initial magic state $|M\rangle^{\otimes n}$?” This size can be computed and it turns out to be small. It implies a new upper bound for the length of the bit strings that is within a factor of 4 of the matching lower bound [11]. This is the content of our main result, Theorem 2, presented in Sec. IV.

II. QUANTUM STATE AS A STRING OF BITS

A crucial feature of the Λ -polytope formalism is that the quantum state $|\Psi(t)\rangle$ of the system is replaced by a bit string $b(t)$ of bounded length. *That bit string $b(t)$ is a valid and accurate representation of the quantum system.*

The role of the bit string $b(t)$ is illustrated in Fig. 1. QCM consists of preparing a quantum register in a magic state $|M\rangle^{\otimes n}$, followed by a sequence of Pauli measurements [see Fig. 1(a)]. This requires a device Prep to deliver the magic states to the quantum register QR and a classical register CR to store the previous measurement record

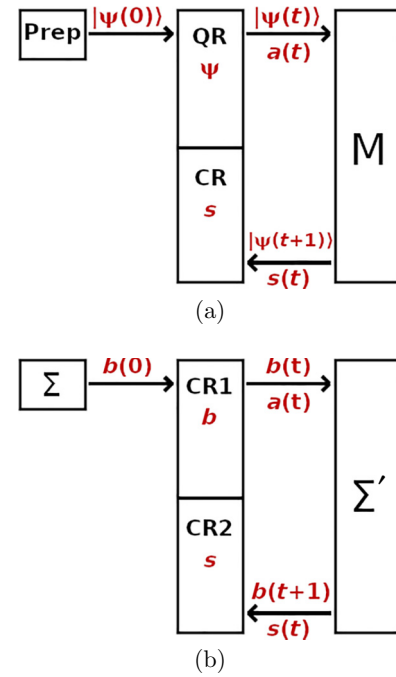


FIG. 1. Quantum computation with magic states (a) and its simulation based on Λ polytopes (b).

s , a classical side computation to identify the label $a(t)$ of the Pauli observable measured in step t , and a measurement device M to perform the measurements and to output the corresponding results $s(t)$. The overall structure of the classical simulation is the same but with the components modified [see Fig. 1(b)]. Prep is replaced by a first sampler Σ that samples from the phase-space distribution of the initial state $|M\rangle^{\otimes n}$. There are two classical registers, CR1 and CR2. The former stores the phase-space samples $b(t)$, and the latter the prior measurement record, as in Fig. 1(a). The measurement device M is replaced by a second sampler Σ' that takes as input a phase-space point $b(t)$ and a Pauli label $a(t)$ and outputs a new phase-space point $b(t+1)$ as well as a measurement outcome $s(t)$. The same information that, in the standard quantum mechanical description, is carried by the quantum state $|\Psi(t)\rangle$ is, in the Λ -polytope description, carried by the bit string $b(t)$.

For universal quantum computation in the magic state model, the classical simulation needs to reproduce the quantum mechanical prediction for the joint distribution of outcomes for any sequence of Pauli measurements. The Λ -polytope method does that, without any approximation. For the predictions, the statistical distribution of the bit strings $\{b(t), \forall t\}$ matters, not individual values $b(t)$. This is the same for the quantum mechanical states $|\Psi(t)\rangle$. They too are conditioned on prior measurement outcomes and hence are probabilistic.

III. PRELIMINARIES

A. Quantum computation with magic states

In Ref. [2], a hidden-variable model (HVM) is defined for quantum computation with magic states (QCM) [13]—a universal model of quantum computation in which computation proceeds through a sequence of Clifford gates and Pauli measurements on an initially prepared “magic state.”

The measurements come from the n -qubit Pauli group \mathcal{P}_n , the group generated by the Pauli operators X , Y , and Z acting on n qubits. Modding out overall phases, we have $\mathcal{P}_n/\mathcal{Z}(\mathcal{P}_n) \cong \mathbb{Z}_2^{2n}$, where $\mathcal{Z}(\mathcal{P}_n) \cong \langle 1, i, -1, -i \rangle$ is the center of the group \mathcal{P}_n and, without loss of generality, we can fix a phase convention for the Pauli operators to be

$$T_a = i^{-\langle a_z | a_x \rangle} \bigotimes_{k=1}^n Z^{a_z[k]} X^{a_x[k]},$$

$\forall a = (a_z, a_x) \in \mathbb{Z}_2^n \times \mathbb{Z}_2^n =: E_n$, where the inner product $\langle a_z | a_x \rangle$ is computed mod 4. The projector corresponding to a Pauli measurement $a \in E_n$ yielding outcome $s \in \mathbb{Z}_2$ is denoted $\Pi_a^s := (1 + (-1)^s T_a)/2$. The gates of the model are chosen from the Clifford group—the normalizer of the Pauli group in the unitary group up to overall phases, $\mathcal{C}\ell_n = \mathcal{N}(\mathcal{P}_n)/U(1)$. They are defined by the property that they map Pauli operators to Pauli operators under conjugation. For more background on QCM, see Appendix B.

B. The Λ -polytope model

We denote by $\text{Herm}(\mathcal{H})$ the space of Hermitian operators on Hilbert space \mathcal{H} and, unless otherwise specified, \mathcal{H} (or \mathcal{H}_n) is the n -qubit Hilbert space $(\mathbb{C}^2)^{\otimes n}$. $\text{Herm}_1(\mathcal{H})$ is the affine subspace of $\text{Herm}(\mathcal{H})$ consisting of operators with unit trace and $\text{Herm}_1^{\geq 0}(\mathcal{H})$ is the subset of $\text{Herm}_1(\mathcal{H})$ consisting of positive semidefinite operators. $\text{Herm}_1^{\geq 0}(\mathcal{H})$ contains the density operators representing physical quantum states.

The state space of the hidden-variable model of Ref. [2] is based on the Λ polytopes. Denoting the set of pure n -qubit stabilizer states by \mathcal{S}_n , the Λ polytope for n qubits is defined as

$$\Lambda_n = \{X \in \text{Herm}_1(\mathcal{H}_n) \mid \text{Tr}(|\sigma\rangle\langle\sigma| X) \geq 0 \forall |\sigma\rangle \in \mathcal{S}_n\}. \quad (1)$$

For any fixed number $n \in \mathbb{N}$ of qubits, Λ_n is a bounded polytope with a finite number of vertices [9]. We denote the vertices of Λ_n by $\{A_\alpha \mid \alpha \in \mathcal{V}_n\}$, where \mathcal{V}_n is an index set [14]. The hidden-variable model is defined by the following theorem.

Theorem 1 (Ref. [2], Theorem 1 and Ref. [9], Theorem 1). For any number of qubits $n \in \mathbb{N}$,

ALGORITHM 1. Classical simulation of a single run of a magic state quantum circuit with input state ρ .

```

1: sample  $\alpha \in \mathcal{V}_n$  according to  $p_\rho : \mathcal{V}_n \rightarrow \mathbb{R}_{\geq 0}$ 
2: propagate  $\alpha$  through the circuit
3: while the end of the circuit has not been reached do
4:   if a Clifford gate  $g \in \mathcal{C}\ell_n$  is encountered then
5:     update the phase space point according to  $\alpha \leftarrow g \cdot \alpha$ 
6:   end if
7:   if a Pauli measurement  $a \in E_n$  is encountered then
8:     sample  $(\beta, s) \in \mathcal{V}_n \times \mathbb{Z}_2$  according to  $q_{\alpha, a}$ 
9:     return  $s \in \mathbb{Z}_2$  as the outcome of the measurement
10:    update the phase space point according to  $\alpha \leftarrow \beta$ 
11:   end if
12: end while

```

(1) Any n -qubit quantum state $\rho \in \text{Herm}_1^{\geq 0}(\mathcal{H}_n)$ can be decomposed as

$$\rho = \sum_{\alpha \in \mathcal{V}_n} p_\rho(\alpha) A_\alpha, \quad (2)$$

with $p_\rho(\alpha) \geq 0$ for all $\alpha \in \mathcal{V}_n$ and $\sum_{\alpha} p_\rho(\alpha) = 1$. That is, any n -qubit quantum state ρ can be represented by a probability distribution p_ρ over \mathcal{V}_n .

(2) For any $A_\alpha, \alpha \in \mathcal{V}_n$ and any Clifford gate $g \in \mathcal{C}\ell_n$, $gA_\alpha g^\dagger$ is a vertex of Λ_n . This defines an action of the Clifford group on \mathcal{V}_n as $gA_\alpha g^\dagger =: A_{g \cdot \alpha}$, where $g \cdot \alpha \in \mathcal{V}_n$.

(3) For any $A_\alpha, \alpha \in \mathcal{V}_n$ and any Pauli projector Π_a^s , we have

$$\Pi_a^s A_\alpha \Pi_a^s = \sum_{\beta \in \mathcal{V}_n} q_{\alpha, a}(\beta, s) A_\beta, \quad (3)$$

with $q_{\alpha, a}(\beta, s) \geq 0$ for all $\beta \in \mathcal{V}_n$ and $s \in \mathbb{Z}_2$, and $\sum_{\beta, s} q_{\alpha, a}(\beta, s) = 1$. That is, Pauli measurements are represented by a stochastic map from (phase-space-point, measurement) pairs $(\alpha, a) \in \mathcal{V}_n \times E_n$ to (phase-space-point, measurement-outcome) pairs $(\beta, s) \in \mathcal{V}_n \times \mathbb{Z}_2$.

A classical simulation algorithm for QCM based on sampling from the defining probability distributions of this HVM is given in Algorithm 1. The algorithm returns samples from the distribution of measurement outcomes for the quantum circuit being simulated that agree with the predictions of quantum theory [2].

IV. MAIN RESULT

In this section, we present our main result—the amount of classical data that the simulation procedure Algorithm 1 must track is small. First, in Sec. IV A, we give the result for a simplified version of the computational model wherein we only allow Pauli measurements, no Clifford

gates, and, in particular, we allow only sequences of independent and commuting Pauli measurements. This simplified model is still universal for quantum computation [15,16]. Then, in Sec. IV B, we give a more general statement of the main result where we allow computations consisting of any sequence of Clifford gates and Pauli measurements.

A. Simplified case

Although the most general quantum computation in QCM could consist of any sequence of Clifford gates and Pauli measurements performed on an arbitrary input state, we can make several assumptions simplifying the computational model while preserving the property of quantum computational universality.

First, in QCM, we usually assume that every computation starts from a fixed magic input state. For example, using the standard magic state circuit gadget [17, Fig. 10.25], any Clifford+ T circuit with n T gates acting on m qubits can be turned into a QCM circuit (Clifford gates and Pauli measurements only) on $n + m$ qubits acting on a state of the form $|0\rangle^{\otimes m} \otimes |H\rangle^{\otimes n}$, where $|H\rangle = (|0\rangle + \exp(i\pi/4)|1\rangle)/\sqrt{2}$. In general, it suffices to consider input states of the form $|0\rangle^{\otimes m} \otimes |M\rangle^{\otimes n}$, where $|M\rangle$ is a fixed single-qubit magic state [18].

Second, we can do away with the Clifford gates altogether [7,15]. To see this, note that the Clifford gates can always be propagated forward in time through the circuit, conjugating the Pauli measurements into other Pauli measurements. Once they are propagated past the final measurement in the circuit, they can be removed, since they no longer affect the statistics of the measurement outcomes.

We can also do away with the stabilizer part of the input. The Pauli circuit on the $(m + n)$ -qubit state $|0\rangle^{\otimes m} \otimes |M\rangle^{\otimes n}$ can be simulated by an adaptive sequence of Pauli measurements acting only on the n -qubit magic part of the input [16,19]. The stabilizer part is handled by extra classical side processing.

Finally, it suffices to consider only sequences of commuting Pauli measurements up to length n [16]. This is because when a Pauli measurement is encountered that anticommutes with a measurement that was previously performed, the outcome of the measurement will be uniformly random and the update of the state after the measurement can be implemented by a Clifford gate. Therefore, the anticommuting measurement can always be replaced by a coin flip to determine the measurement outcome and a Clifford gate to implement the state update, which can then be propagated past the future measurements. The longest sequences of independent pair-wise commuting Pauli measurements on n qubits have length n .

ALGORITHM 2. Classical simulation of a single run of a Pauli-based quantum circuit with input state ρ .

```

1: sample  $\alpha_0 \in \mathcal{V}_n$  according to  $p_\rho : \mathcal{V}_n \rightarrow \mathbb{R}_{\geq 0}$ 
2: for all  $a_t, t \in \{1, 2, \dots, n\}$  do
3:   sample  $(\alpha_t, s_t) \in \mathcal{V}_n \times \mathbb{Z}_2$  according to  $q_{\alpha_{t-1}, \alpha_t}$ 
4:   return  $s_t \in \mathbb{Z}_2$  as the outcome of measurement  $a_t$ 
5: end for

```

To summarize, for universal quantum computation, it suffices to consider adaptive sequences of pair-wise commuting Pauli measurements of length n acting on a fixed magic state of the form $|M\rangle^{\otimes n}$. A simplified version of the classical simulation algorithm above for circuits of this form is given by Algorithm 2. For a complete description of the above circuit simplifications, along with a method for compiling a given QCM circuit into an adaptive Pauli circuit, see Ref. [16].

We can now state our main result.

Theorem 2 (Main result). Any quantum computation consisting of a sequence of n independent pair-wise commuting Pauli measurements on a fixed magic state $|M\rangle^{\otimes n}$ can be simulated by Algorithm 2 using a memory of $2n^2 + 3n$ bits to specify the phase-space points reached.

Note that the precise form of the magic input state used is not important for the proof of this theorem, so long as the state is fixed for each number of qubits. For universal quantum computation, it suffices to consider n copies of a fixed single-qubit magic state such as $|H\rangle^{\otimes n}$. The proof of this theorem relies on the following result from convex geometry (see, e.g., Ref. [20, Sec. 1.6]).

Theorem 3 (Carathéodory's theorem). If a point x of \mathbb{R}^D lies in the convex hull of a set V , then x can be written as the convex combination of at most $D + 1$ points in V .

Proof of Theorem 2. Since the generalized phase-space-point operators $\{A_\alpha | \alpha \in \mathcal{V}_n\}$ are not a basis for $\text{Herm}_1(\mathcal{H}_n)$, they are overcomplete; the distributions p_ρ that represent states and the distributions $q_{\alpha,a}$ that represent Pauli measurements in the model of Theorem 1 are not unique. Since Λ_n lives in $\text{Herm}_1(\mathcal{H}_n)$, a real affine space of dimension $4^n - 1$, by Carathéodory's theorem there exist choices for the distribution p_ρ such that $|\text{supp}(p_\rho)| \leq 4^n$. Similarly, for each $s \in \mathbb{Z}_2$, there exist choices for $q_{\alpha,a}(-, s)$ such that $|\text{supp}(q_{\alpha,a}(-, s))| \leq 4^n$. To start, we fix a canonical choice of the probability distributions p_ρ and $q_{\alpha,a}$ satisfying these properties. With the distribution p_ρ representing the input state of the circuit fixed, we can label the elements of the support by the numbers $1, 2, \dots, 4^n$, and so specifying a sample from this distribution (i.e., specifying an element of the support of p_ρ) amounts to specifying




FIG. 2. Three descriptions of a quantum computation. (a) A circuit-level description. Pauli measurements a_1, a_2, \dots are performed on the input state ρ , yielding measurement outcomes s_1, s_2, \dots (b) The representation of this computation in the model of Theorem 1. The input state is represented by the probability distribution p_ρ , each measurement a_t being represented by a probability distribution q_{α_{t-1}, a_t} . (c) A single run of the simulation algorithm based on the probabilistic representation of the computation. We start by sampling from p_ρ to obtain α_0 . For each measurement a_t , we sample from q_{α_{t-1}, a_t} to obtain measurement outcome s_t and updated state α_t . (d) An upper bound on the number of classical bits required to store each piece of the sampling record of (c) according to Theorem 2.

an integer between 1 and 4^n , which takes no more than $\log_2(4^n) = 2n$ bits.

There are $4^n - 1$ nontrivial n -qubit Pauli measurements; therefore, specifying each measurement requires no more than $2n$ bits. For the t th measurement a_t , the distribution q_{α_{t-1}, a_t} is uniquely specified by the sampling record consisting of states $\alpha_0, \alpha_1, \dots, \alpha_{t-1}$, measurements a_1, a_2, \dots, a_{t-1} , and measurement outcomes s_1, s_2, \dots, s_{t-1} . Once this distribution is fixed, with the canonical choice above, specifying a sample from this distribution requires no more than $2n + 1$ bits (1 bit for s_t , and as before $2n$ bits for α_t).

Since the length of the measurement sequence is no more than n , the number of classical bits required to specify the complete sampling record is no more than

$$\underbrace{2n}_{\alpha_0} + \sum_{t=1}^n \left[\underbrace{2n}_{a_t} + \underbrace{1}_{s_t} + \underbrace{2n}_{\alpha_t} \right] = 4n^2 + 3n.$$

For a summary of the accounting used to obtain this bound, see Fig. 2.

This initial bound can be improved in a number of ways. First, for the purpose of simulation, we do not need to store α_n , since there are no more measurements. This immediately removes $2n$ bits. Second, after the t th measurement, the value of any measurement in the span of a_1, a_2, \dots, a_t is already determined. Therefore, if the $(t+1)$ th measurement a_{t+1} is to be independent and commute with the previous measurements, it is chosen from

$\{a_1, a_2, \dots, a_t\}^\perp / \text{span}(a_1, a_2, \dots, a_t) \cong E_{n-t}$. Specifying a measurement chosen from this set requires only $2(n-t)$ bits, not the full $2n$ bits.

Finally, we can perform another simplification that reduces the number of qubits by 1 after each measurement. After a measurement of $a \in E_n$ giving outcome $s \in \mathbb{Z}_2$, the relevant state space is projected to $\Pi_a^s \Lambda_n \Pi_a^s$. This is contained in a $(4^{n-1} - 1)$ -dimensional subspace of $\text{Herm}_1(\mathcal{H}_n)$. There exists a Clifford gate $g \in \mathcal{C}\ell_n$ such that $g \Pi_{z_n}^0 g^\dagger = \Pi_a^s$, and $\Pi_a^s \Lambda_n \Pi_a^s = g \Pi_{z_n}^0 \Lambda_n \Pi_{z_n}^0 g^\dagger = g(\Lambda_{n-1} \otimes |0\rangle\langle 0|)g^\dagger$. Therefore, after the measurement, the Clifford gate g can be propagated out, yielding a Pauli-measurement circuit on an input state with a stabilizer state tensor factor. Then the above-mentioned circuit simplifications [16,19] can be used to remove the stabilizer part of the input. This reduction can be performed after each measurement. This is similar to the idea behind the reduced classical simulation of Theorem 3 of Ref. [19]. With this dimension reduction after each measurement, to specify the sample (α_t, s_t) requires no more than $2(n-t) + 1$ bits. This dimension reduction implicitly implements both of the reductions of the previous paragraph. Note that this explicit dimension reduction is not required to obtain the memory savings since, according to Ref. [19, Theorem 2], $\Lambda_{n-t} \otimes |0\rangle\langle 0|^{\otimes t}$ is a subpolytope of Λ_n contained in a $(4^{n-t} - 1)$ -dimensional affine subspace of $\text{Herm}_1(\mathcal{H}_n)$ and so by Carathéodory's theorem, an operator of the form $X \otimes |0\rangle\langle 0|^{\otimes t}$ can be decomposed into vertices of this subpolytope with support of size no more than 4^{n-t} .

With these reductions, the complete measurement record can be specified with no more than

$$\underbrace{2n}_{\alpha_0} + \sum_{t=1}^n \left[\underbrace{2(n-t+1)}_{a_t} + \underbrace{1}_{s_t} + \underbrace{2(n-t)}_{\alpha_t} \right] = 2n^2 + 3n$$

classical bits, which is the claimed bound. \blacksquare

An idea similar to the one used in the proof of Theorem 2, of explicitly tracking the measurement record of a quantum circuit, has previously been used to define a contextual hidden-variable model for the stabilizer subtheory [21].

B. General case

Theorem 2 uses a simplified version of QCM, where, without loss in computational power, all Clifford gates have been eliminated and the measurement sequence is shrunk to at most n commuting Pauli measurements. In the remainder of this section, we demonstrate that the bound on the size of the reachable phase-space region does not increase much if we do not make these simplifications, i.e., if we admit arbitrarily long sequences of (potentially noncommuting) Pauli measurements, and Clifford gates between them.

Specifically, we have the following additional result.

Corollary 1. Any quantum computation consisting of an arbitrarily long sequence of Pauli measurements and Clifford unitaries, applied to a fixed magic state $|M\rangle^{\otimes n}$, can be simulated using a memory of $4n^2 + 6n$ bits to specify the reachable phase-space points.

Comparing with Theorem 2, we find that the memory requirement merely doubles; in particular, the quadratic scaling with the number n of magic states remains unchanged.

Proof of Corollary 1. Every vertex of Λ_n is in exactly one orbit of vertices with respect to the action of the Clifford group. Therefore, the number of Clifford orbits traveled by Algorithm 2 is smaller than or at most equal to the number of vertices traveled. A consequence of Theorem 2 thus is that the number of equivalence classes traveled by Algorithm 2 is no more than 2^{2n^2+3n} .

At every given point in the classical simulation, the single vertex A_v under consideration by the simulation algorithm has a product structure, $A_v = g(\tilde{A}_v \otimes |S\rangle\langle S|)g^\dagger$, where $|S\rangle\langle S|$ is the projector onto an m -qubit stabilizer state with stabilizer group S , \tilde{A}_v is a vertex of Λ_{n-m} , and g is a Clifford unitary [19].

When switching from Algorithm 2 specialized to the simplified computational model back to the general Algorithm 1, the following additional situations may occur at any given step: (i) the next operation is a Clifford unitary, (ii) the next operation is the measurement of a Pauli observable T_a , $a \in E_n$, such that $\pm T_a \in gSg^\dagger$, and (iii) the next operation is the measurement of a Pauli observable T_a that does not commute with all elements of gSg^\dagger .

In case (i), the Clifford orbit $[A_v]$ of A_v does not change. In case (ii), the measurement outcome deterministically follows from the stabilizer and no update of A_v occurs at all. Hence, no update of $[A_v]$ occurs either. In case (iii), the measurement outcome is uniformly random and for either outcome, the ensuing transformation can be replaced by a Clifford unitary. Thus we are back to case (i)—no update of $[A_v]$ occurs. The conclusion is that the number of Clifford orbits of vertices reachable by the more general Algorithm 1 equals the number of orbits reachable by the specialized Algorithm 2 applying to the canonical form of QCMs.

Now, the size of every Clifford orbit of vertices of Λ_n is upper bounded by the size $|\mathcal{C}\ell_n|$ of the n -qubit Clifford group. Therefore, the number of vertices reachable by Algorithm 1 is bounded by $2^{2n^2+3n} \times |\mathcal{C}\ell_n|$. The number of bits required to specify an n -qubit Clifford gate is no more than $2n^2 + 3n$. Namely, the conjugation action of a Clifford unitary on a Pauli operator is described by a binary-valued $2n \times 2n$ matrix plus a $2n$ -component binary vector for the signs of Pauli operators [17], requiring $2n \times (2n + 1)$ bits.

From this, the $2n^2 - n$ constraints imposed by preservation of the Pauli commutation relations need to be subtracted. Hence, to specify a vertex reachable in the simulation using Algorithm 1, we specify a Clifford orbit $[A_v]$ and a Clifford operation g , which takes no more than $4n^2 + 6n$ bits, as claimed. ■

We may generalize beyond the scenario of Corollary 1 by supplementing the initial magic state $|H\rangle^{\otimes n}$ with an m -qubit ancilla in a stabilizer state, where the stabilizer part can be kept separate and tracked independently in the simulation. As in the previous generalization, the computational power does not increase but a broader physical setting is represented.

Then, the number of bits needed depends on both the number of qubits m supporting the stabilizer part of the input and the number of magic states n , namely,

$$[2m^2 + m] + [4n^2 + 6n] + [2(n + m)^2 + 3(n + m)]. \quad (4)$$

Here, the first term accounts for the stabilizer part of the state, the second term for the magic part (this is the first version of the bound given in the proof of Theorem 2 shown in Fig. 2), and the third term is for the $(m + n)$ -qubit Clifford group element that needs to be tracked, as in the proof of Corollary 1.

In deriving Eq. (4), we make use of the fact that a tensor product of a projector onto an m -qubit stabilizer state with a vertex of Λ_n is a vertex of Λ_{m+n} [19, Theorem 2], and we use the stabilizer-formalism techniques from Ref. [19, Sec. 4] to isolate the update of the stabilizer part from the magic part of the state. The number of bits tracked remains quadratic in n and m .

V. DISCUSSION AND CONCLUSIONS

To summarize, in this work we have shown that the classical simulation of universal quantum computation using the Λ polytopes [2] is a small-data problem. Specifically, it is shown that with respect to the model of quantum computation with magic states, the number of bits that represent the quantum system at any stage of the simulation is quadratic in the number n of the magic states. Classical simulation of quantum computation is (presumably) still hard. In the present case, this hardness does not stem from shuffling around lots of data but, instead, from complicated operations on little data.

For illustration, we compare the above classical simulation method to two others, one very different and the other rather similar. The first method is the straightforward simulation of a quantum system obtained by choosing a specific basis of the Hilbert space at hand, mapping operators to matrices and states to vectors. The state of the system at any moment in time is now described by exponentially many complex-valued amplitudes. The bit equivalent of a

complex number is a matter of numerical precision but in any case the amount of data to be processed is large. The state update is conceptually simple and the computational hardness derives from the size of the objects involved. This simulation method and that of Ref. [2] thus represent opposite ends of the spectrum.

The second simulation method with which we compare is that of sampling from Wigner functions [4], applicable to odd Hilbert-space dimension. Here, the overall structure of the simulation is the same as in Ref. [2], i.e., repeated sampling from a phase space. There are, however, two important differences. In Ref. [4], (i) the sampling is computationally efficient whenever it applies but (ii) the sampling procedure does not apply to all initial magic states. Specifically, it only applies when the Wigner function [5,6] of the initial state is positive. Indeed, this is why negativity of the Wigner function is a precondition for computational speed-up. In Ref. [2], (i) the sampling is not guaranteed to be computationally efficient but (ii) it applies to all possible initial states.

We highlight two further aspects of the simulation method [2]:

- (a) In the classical simulation of quantum computation using Λ polytopes, the description of the state of the system by $2n^2 + 3n$ bits does not invoke any approximation. The distributions of measurement outcomes sampled from are the exact quantum mechanical ones. Thus, the data representing the system are genuinely discrete [22] and for this reason we regard Ref. [2] as a realization of Wheeler's "it from bit" proposal.
- (b) Theorem 2 can be generalized to the odd-prime-dimensional qudit case [9]. In this case, the precompilation step reducing the computation to sequences of commuting Pauli measurements is also possible [23] but again not necessary.

The Λ polytopes are only beginning to be explored and they may hold many more surprises. From the perspective of quantum computation, an important question is the following. Two known efficient classical simulation methods, the stabilizer formalism [24] and sampling from non-negative Wigner functions [4], are reproduced and extended by the Λ -polytope method, at the level of the simplest vertices of the polytopes. Where does efficient classical simulation end?

ACKNOWLEDGMENTS

M.Z. and R.R. thank William G. Unruh for insightful discussions. M.Z. and R.R. are funded by the Natural Sciences and Engineering Research Council of Canada (NSERC), in part through the Canada First Research

Excellence Fund, Quantum Materials and Future Technologies Program. R.R. is supported by the Alexander von Humboldt Foundation. C.O. is supported by the U.S. Air Force Office of Scientific Research under Award No. FA9550-21-1-0002. This work was also supported by the Horizon Europe project FoQaCiA, Grant Agreement No. 101070558.

APPENDIX A: UPPER BOUND ON THE NUMBER OF VERTICES OF Λ

Lemma 1. For any number of qubits $n \in \mathbb{N}$, the number of vertices $|\mathcal{V}_n|$ of Λ_n satisfies

$$\log_2(|\mathcal{V}_n|) \leq 1 + 4^{n-1} [n^2 - n + 2 \log_2(10e)].$$

Proof of Lemma 1. According to the upper-bound theorem of polytope theory [12] (for a more detailed exposition on the upper-bound theorem, see Ref. [20, Sec. 8.4]), the number of facets of a D -dimensional polytope with v vertices is bounded by the number of facets of the D -dimensional cyclic polytope with v vertices, denoted $C(v, D)$. By duality, the number of vertices of the Λ polytope on n qubits is bounded by the number of vertices of the polar dual of the $4^n - 1$ -dimensional cyclic polytope with $|\mathcal{S}_n|$ vertices.

The number of vertices of the dual of $C(v, D)$ (equivalently, the number of facets of $C(v, D)$) is [29, Sec. 4.7],

$$f_{d-1}(C(v, D)) = \begin{cases} \frac{v}{v-m} \binom{v-m}{m}, & \text{for even } D = 2m, \\ 2 \binom{v-m-1}{m}, & \text{for odd } D = 2m + 1. \end{cases}$$

In the case of Λ_n , $D = 4^n - 1$, i.e., D is odd and $m = 2^{2n-1} - 1$. The number of facets of Λ_n is $|\mathcal{S}_n|$. Therefore,

$$\begin{aligned} |\mathcal{V}_n| &\leq f_{4^n-2}(C(|\mathcal{S}_n|, 4^n - 1)) \\ &= 2 \binom{|\mathcal{S}_n| - 2^{2n-1}}{2^{2n-1} - 1} \\ &\leq 2 \binom{|\mathcal{S}_n| - 2^{2n-1}}{2^{2n-1}} \\ &\leq 2 \left(\frac{e(|\mathcal{S}_n| - 2^{2n-1})}{2^{2n-1}} \right)^{2^{2n-1}}. \end{aligned}$$

In the last line, we use a standard upper bound for the binomial coefficient $\binom{n}{k} \leq \left(\frac{en}{k}\right)^k$. Then,

$$\begin{aligned} \log_2(|\mathcal{V}_n|) &\leq 1 + 2^{2n-1} [\log_2(|\mathcal{S}_n| - 2^{2n-1}) \\ &\quad - 2n + 1 + \log_2(e)] \\ &\leq 1 + 2^{2n-1} \left[\log_2\left(5 \cdot 2^{(n^2+3n)/2}\right) \right. \\ &\quad \left. - 2n + 1 + \log_2(e) \right] \\ &\leq 1 + 2^{2n-1} \left[\frac{n^2 + 3n}{2} - 2n + 1 + \log_2(5e) \right] \\ &= 1 + 4^{n-1} [n^2 - n + 2 \log_2(10e)]. \end{aligned}$$

Here, in the second line, we use the upper bound on the number of n -qubit stabilizer states, $|\mathcal{S}_n| \leq 5 \cdot 2^{(n^2+3n)/2}$, from Ref. [30]. This proves the upper bound. ■

We also have the following somewhat simpler bound.

Corollary 2. For any number of qubits $n \in \mathbb{N}$, $\log_2(|\mathcal{V}_n|) \leq 4^n n^2$.

Proof of Corollary 2. For $n = 1$ and $n = 2$, we can enumerate the vertices of Λ_n and we find that the numbers of vertices are 8 and 22320, respectively. These both satisfy the bound. Further, for $n \geq 2$, $1 + 4^{n-1}[n^2 - n + 2 \log_2(10e)] \leq 4^n n^2$ and so the remaining cases follow immediately from Lemma 1. ■

APPENDIX B: QUANTUM COMPUTATION WITH MAGIC STATES

QCM is a universal model of quantum computation in which the stabilizer subtheory (stabilizer states, Clifford gates, and Pauli measurements) is supplemented with so-called “magic” (nonstabilizer) states to achieve quantum computational universality [13]. The stabilizer subtheory is not universal for quantum computation and circuits consisting of elements from this subtheory can be efficiently simulated by a classical computer [24,31].

In the circuit model, universality is provided by the capacity to perform non-Clifford unitary gates. In the magic state model, on the other hand, universality is restored by the inclusion of nonstabilizer input states. In Fig. 3, we show a standard example of this phenomenon: a non-Clifford T gate is performed, using only stabilizer operations, by injecting the magic state $|H\rangle = (|0\rangle + e^{i\pi/4}|1\rangle)/\sqrt{2}$. Therefore, given access to operations in the stabilizer subtheory, as well as copies of a suitable magic state such as $|H\rangle$, any quantum computation can be performed.

For quantum computational universality, it suffices to supplement the stabilizer operations with copies of a




FIG. 3. A quantum circuit showing the implementation of the non-Clifford T gate in the QCM framework, via the injection of the magic state $|H\rangle = (|0\rangle + e^{i\pi/4}|1\rangle)/\sqrt{2}$. Since the Clifford+ T gate set is universal for quantum computation, this proves that, with access to suitable magic states, QCM is also universal. Circuit reproduced from Ref. [17, Fig. 10.25].

single-qudit magic state such as $|H\rangle$. Of course, without the capacity to perform nonstabilizer operations, it is not clear that these states can be prepared. Fortunately, there exist protocols called magic state distillation protocols in which multiple noisy copies of a magic state are consumed to produce fewer copies of the magic state with higher fidelity. This approach is one of the leading candidates for scalable fault-tolerant quantum computation [32].

For more on quantum computation with magic states, see, e.g., Ref. [32, Sec. II] or Ref. [17, Ch. 10] or, for a more detailed summary, see Ref. [2, Sec. SM.II].

-
- [1] J. A. Wheeler, in *Proceedings of the 3rd International Symposium on Foundations of Quantum Mechanics in the Light of New Technology*, edited by H. Ezawa, S. I. Kobayashi, and Y. Murayama (Physical Society of Japan, Tokyo, 1989), p. 354 (Also published in Ref. [25]).
 - [2] M. Zurel, C. Okay, and R. Raussendorf, Hidden variable model for universal quantum computation with magic states on qubits, *Phys. Rev. Lett.* **125**, 260404 (2020).
 - [3] E. F. Galvão, Discrete Wigner functions and quantum computational speedup, *Phys. Rev. A* **71**, 042302 (2005).
 - [4] V. Veitch, C. Ferrie, D. Gross, and J. Emerson, Negative quasi-probability as a resource for quantum computation, *New J. Phys.* **14**, 113011 (2012).
 - [5] D. Gross, Hudson’s theorem for finite-dimensional quantum systems, *J. Math. Phys.* **47**, 122107 (2006).
 - [6] D. Gross, Ph.D. thesis, Institute for Mathematical Sciences, Imperial College London, 2008, <https://www.thp.uni-koeln.de/gross/files/diss.pdf>.
 - [7] N. Delfosse, P. Allard Guerin, J. Bian, and R. Raussendorf, Wigner function negativity and contextuality in quantum computation on rebits, *Phys. Rev. X* **5**, 021003 (2015).
 - [8] R. Raussendorf, J. Bermejo-Vega, E. Tyhurst, C. Okay, and M. Zurel, Phase-space-simulation method for quantum computation with magic states on qubits, *Phys. Rev. A* **101**, 012350 (2020).
 - [9] M. Zurel, C. Okay, R. Raussendorf, and A. Heimendahl, Hidden variable model for quantum computation with magic states on qudits of any dimension, *Quantum* **8**, 1323 (2024).
 - [10] Elsewhere in the quantum domain, there are sampling problems that are algorithmically hard classically but for which samples can be obtained quantum mechanically; e.g., boson

- sampling [26], random circuit sampling [27], and certain instantaneous quantum circuits [28].
- [11] A. Karanjai, J. J. Wallman, and S. D. Bartlett, Contextuality bounds the efficiency of classical simulation of quantum processes, [arXiv:1802.07744](https://arxiv.org/abs/1802.07744).
- [12] P. McMullen, The maximum numbers of faces of a convex polytope, *Mathematika* **17**, 179 (1970).
- [13] S. Bravyi and A. Kitaev, Universal quantum computation with ideal Clifford gates and noisy ancillas, *Phys. Rev. A* **71**, 022316 (2005).
- [14] Λ_n refers to a specific mathematical structure, a polytope, for a fixed number of qubits n . We refer to these polytopes collectively as the Λ polytopes, or Λ (no subscript). The set of vertices of Λ_n is $\{A_\alpha \mid \alpha \in \mathcal{V}_n\}$. That is, A_α is a vertex of Λ_n and it is labeled by $\alpha \in \mathcal{V}_n$. \mathcal{V}_n is the set of all labels of vertices of Λ_n . There is no n subscript on A_α (or α) even though it refers to an object associated with a fixed number of qubits n , because this fact is implicit in the statement $\alpha \in \mathcal{V}_n$.
- [15] S. Bravyi, G. Smith, and J. A. Smolin, Trading classical and quantum computational resources, *Phys. Rev. X* **6**, 021043 (2016).
- [16] F. C. R. Peres and E. F. Galvão, Quantum circuit compilation and hybrid computation using Pauli-based computation, *Quantum* **7**, 1126 (2023).
- [17] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, Cambridge, UK, 2010).
- [18] B. W. Reichardt, Quantum universality by state distillation, *Quantum Inf. Comput.* **9**, 1030 (2009).
- [19] C. Okay, M. Zurek, and R. Raussendorf, On the extremal points of the Λ -polytopes and classical simulation of quantum computation with magic states, *Quantum Inf. Comput.* **21**, 1091 (2021).
- [20] G. M. Ziegler, *Lectures on Polytopes* (Springer Science+Business Media, New York, 1995).
- [21] A. Karanjai, Ph.D. thesis, School of Physics, University of Sydney, 2019, <https://hdl.handle.net/2123/22134>.
- [22] Any continuous parameter dependence is relegated to the sampling probabilities in the first sampling step.
- [23] F. C. R. Peres, Pauli-based model of quantum computation with higher-dimensional systems, *Phys. Rev. A* **108**, 032606 (2023).
- [24] D. Gottesman, in *Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics*, edited by S. P. Corney, R. Delbourgo, and P. D. Jarvis (International Press of Boston, Inc., 1999), p. 32.
- [25] J. A. Wheeler, in *Complexity, Entropy and the Physics of Information*, edited by W. H. Zurek, The Proceedings of the 1988 Workshop on Complexity, Entropy, and the Physics of Information, held May-June, 1989 in Santa Fe, New Mexico (CRC Press, LLC, Santa Fe, New Mexico, 1990), 1st ed., p. 3.
- [26] S. Aaronson and A. Arkhipov, in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (2011), p. 333.
- [27] A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani, On the complexity and verification of quantum random circuit sampling, *Nat. Phys.* **15**, 159 (2019).
- [28] D. Shepherd and M. J. Bremner, Temporally unstructured quantum computation, *Proc. R. Soc. A* **465**, 1413 (2009).
- [29] B. Grünbaum, *Convex Polytopes* (Springer, New York, 2003).
- [30] T. Singal, C. Chiang, E. Hsu, E. Kim, H.-S. Goan, and M.-H. Hsieh, Counting stabiliser codes for arbitrary dimension, *Quantum* **7**, 1048 (2023).
- [31] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, *Phys. Rev. A* **70**, 052328 (2004).
- [32] E. T. Campbell, B. M. Terhal, and C. Vuillot, Roads towards fault-tolerant universal quantum computation, *Nature* **549**, 172 (2017).