

Negative quasi-probability as a resource for quantum computation

This content has been downloaded from IOPscience. Please scroll down to see the full text.

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 138.25.78.25

This content was downloaded on 20/08/2017 at 10:39

Please note that [terms and conditions apply](#).

Negative quasi-probability as a resource for quantum computation

This article has been downloaded from IOPscience. Please scroll down to see the full text article.

2012 New J. Phys. 14 113011

(<http://iopscience.iop.org/1367-2630/14/11/113011>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 89.202.245.164

The article was downloaded on 20/03/2013 at 16:31

Please note that [terms and conditions apply](#).

Corrigendum: Negative quasi-probability as a resource for quantum computation

2012 *New J. Phys.* **14** 113011

**Victor Veitch¹, Christopher Ferrie¹, David Gross²
and Joseph Emerson¹**

¹ Institute for Quantum Computing and Department of Applied Mathematics,
University of Waterloo, Waterloo, Ontario N2L 3G1, Canada

² Institute for Physics, University of Freiburg, Rheinstrasse 10,
D-79104 Freiburg, Germany

E-mail: vveitch@uwaterloo.ca, csferrie@uwaterloo.ca,
david.gross@physik.uni-freiburg.de and jemerson@uwaterloo.ca

New Journal of Physics **15** (2013) 039502 (3pp)

Received 4 February 2013

Published 21 March 2013

Online at <http://www.njp.org/>

doi:10.1088/1367-2630/15/3/039502

One entry of each Wigner function appearing in figures 3 and 4 is mislabeled in Victor Veitch *et al* (2012 *New J. Phys.* **14** 113011). The corrected labeling appears in the figures below. The results and conclusions of the paper are unchanged.



Content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](https://creativecommons.org/licenses/by/3.0/).
Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

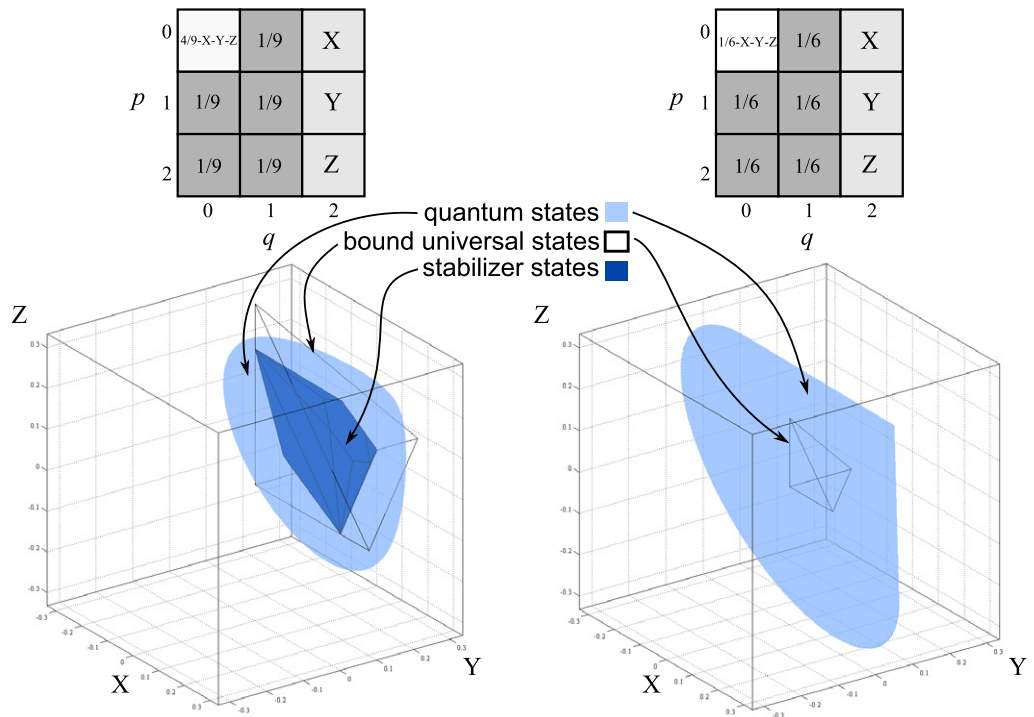


Figure 3. Orthogonal 2D slice of qutrit state space. On the left are the entries of the Wigner function which define the 2D slice that is plotted on the right. Six entries are fixed at $1/9$, while X and Y identify the entries that are free to vary, and the remaining entry is determined by normalization. The maximally mixed state is the point $(X, Y) = (1, 1)/9$. The various regions carved out by varying these values are shown on the right. There are $\binom{8}{2} = 28$ such slices which are identical (up to a relabeling of the axes). These would be the only slices featuring the maximally mixed state.

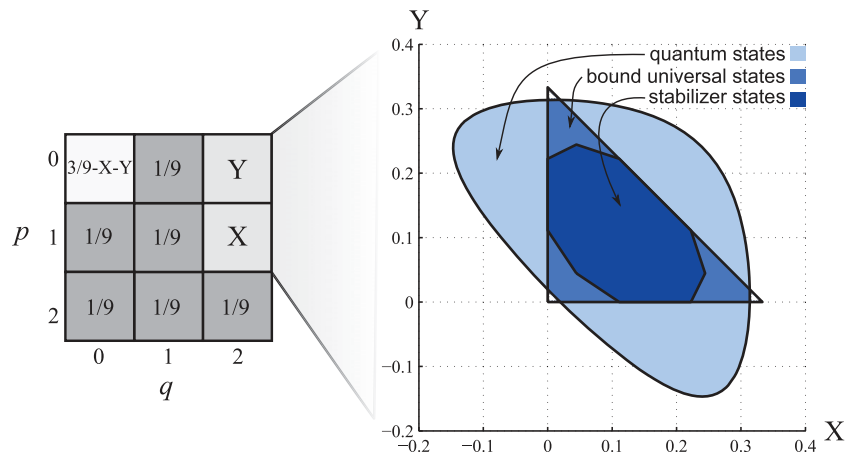


Figure 4. Orthogonal three-dimensional (3D) slices of qutrit state space. Above each figure are the entries of the Wigner function which define the associated 3D slice. Five values are fixed at a value of $1/9$ (left figure) and $1/6$ (right figure), while X , Y and Z identify the entries that are free to vary, and the remaining entry is determined by normalization. In each case, due to symmetries, there are $\binom{8}{3} = 56$ such slices which are identical (up to a relabeling of the axes). Note that the slice on the right does not cut through the stabilizer polytope but does contain a region of bound states. Also note that this slice contains one of the nine states with a maximal negativity of $1/3$ while the slice on the left, and those equivalent up to permutations, are the only ones which feature the maximally mixed state $(X, Y, Z) = 1/9$. See also figure 2 for two-dimensional (2D) slice of the figure on the left.

Acknowledgments

We thank Thom Bohdanowicz for pointing out the error in the original figures.

Negative quasi-probability as a resource for quantum computation

Victor Veitch¹, Christopher Ferrie¹, David Gross²
and Joseph Emerson¹

¹ Institute for Quantum Computing and Department of Applied Mathematics,
University of Waterloo, Waterloo, ON N2L 3G1, Canada

² Institute for Physics, University of Freiburg, Rheinstrasse 10, D-79104
Freiburg, Germany

E-mail: vveitch@uwaterloo.ca, csferrie@uwaterloo.ca,
david.gross@physik.uni-freiburg.de and jemerson@uwaterloo.ca

New Journal of Physics **14** (2012) 113011 (21pp)

Received 22 August 2012

Published 8 November 2012

Online at <http://www.njp.org/>

doi:10.1088/1367-2630/14/11/113011

Abstract. A central problem in quantum information is to determine the minimal physical resources that are required for quantum computational speed-up and, in particular, for fault-tolerant quantum computation. We establish a remarkable connection between the potential for quantum speed-up and the onset of negative values in a distinguished quasi-probability representation, a discrete analogue of the Wigner function for quantum systems of odd dimension. This connection allows us to resolve an open question on the existence of bound states for magic state distillation: we prove that there exist mixed states outside the convex hull of stabilizer states that cannot be distilled to non-stabilizer target states using stabilizer operations. We also provide an efficient simulation protocol for Clifford circuits that extends to a large class of mixed states, including bound universal states.



Content from this work may be used under the terms of the [Creative Commons Attribution-NonCommercial-ShareAlike 3.0 licence](https://creativecommons.org/licenses/by-nc-sa/3.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Contents

1. Introduction	2
2. Previous work	4
3. Background	5
3.1. The stabilizer formalism	5
3.2. Magic state distillation	6
3.3. Quasi-probability representations	7
3.4. The discrete Wigner function	8
4. Negative discrete Wigner representation is necessary for computational speed-up	9
5. Magic state distillation	12
6. The geometry of positively represented states	16
7. Discussion and conclusions	17
Acknowledgments	19
References	19

1. Introduction

While it is widely believed that quantum computers can solve certain problems with exponentially fewer resources than their classical counterparts, the scope of the physical resources of the underlying quantum systems that enable universal quantum computation is not well understood. For example, for the standard circuit model of quantum computation, Vidal has shown that high entanglement is necessary for an exponential speed-up (Vidal 2003); however, it is also known that access to high entanglement is not sufficient (Gottesman 1997). Moreover, in alternative models of quantum computation such as DQC1 (Knill and Laflamme 1998), algorithms that may be performed on highly mixed input states appear to be more powerful than classical computation even though there appears to be a negligible amount of entanglement in the underlying quantum system (Datta *et al* 2005). This suggests that large amounts of entanglement, purity or even coherence may not be necessary resources for quantum-computational speed-up. One of the central open problems of quantum information is to understand which sets of quantum resources are jointly necessary and sufficient to enable an exponential speed-up over classical computation. Any solution to this important problem may point to a more practical experimental means of achieving the benefits of quantum computation.

The question of whether a restricted subset of quantum theory is still sufficient for a given task is meaningful when there is a specific context that divides the full set of possible quantum operations into two classes: the restricted subset of operations that are accessible or easy to implement and the remainder that are not. In such a context it is then natural to consider the difficult operations as resources and ask how much, if any, of these resources are required. For example, a common paradigm in quantum communication is that of two or more spatially separated parties for which local quantum operations and classical communication define a restricted set of operations that are accessible or ‘free resources’, whereas joint quantum operations are not free; in this context entanglement is the natural resource for quantum communication. Here we are interested not in quantum communication but in the power of quantum computation, and in particular the practically relevant case of fault-tolerant quantum

computation. Transversal unitary gates, i.e. gates that do not spread errors within each code block, play a critical role in fault-tolerant quantum computation. Recent theoretical work has shown that a set of quantum gates which is both universal and transversal, and hence fault-tolerant (Zeng *et al* 2007, Chen *et al* 2008, Eastin and Knill 2009), does not exist. That is, any scheme for fault-tolerant quantum computation divides quantum operations into two classes: those with a fault-tolerant implementation—these are the ‘free resources’—and the remainder—these are not free but are required to achieve universality. For a fixed fault-tolerant scheme the critical question is: what are the necessary and sufficient physical resources to promote fault-tolerant computation to universal quantum computation?

Most of the best-known fault-tolerant schemes are built around the well-known stabilizer formalism (Gottesman 2006), in which a distinguished set of preparations, measurements and unitary transformations (the ‘stabilizer operations’) have a fault-tolerant implementation. Stabilizer operations also arise naturally in some physical systems with topological order (Moore and Read 1991, Douçot and Vidal 2002, Lloyd 2002). As described above, the transversal set of stabilizer operations do not give a universal gate set and must be supplemented with some additional (non-stabilizer) resource. A celebrated scheme for overcoming this limitation is the magic state model of quantum computation devised by Bravyi and Kitaev (2004) where the additional resource is a set of ancilla systems prepared in some (generally noisy) non-stabilizer quantum state. Hence, in this important paradigm, the question of which physical resources are required for universal fault-tolerant quantum computation reduces to the following: which non-stabilizer states are necessary and sufficient to promote stabilizer computation to universal quantum computation?

In this paper, we identify a non-trivial closed, convex subset of the space of quantum states which we prove is incapable of producing universal fault-tolerant quantum computation. In particular, we show that this convex subset strictly contains the convex hull of stabilizer states, and thereby prove that there exists a class of *bound universal states*, i.e. states that cannot be prepared from convex combinations of stabilizer states and yet are not useful for quantum computation. Thus our proof of the existence of bound universal states resolves in the negative the open problem raised by Bravyi and Kitaev (2004) of whether *all* non-stabilizer states promote stabilizer computation to universal quantum computation. Furthermore, we devise an efficient simulation algorithm for the subset of quantum theory that consists of operations from the stabilizer formalism acting on inputs from our non-universal region, which includes mixed states both inside and outside the convex hull of stabilizer states. This simulation scheme is an extension of the celebrated Gottesman–Knill theorem (Gottesman 1997, Aaronson and Gottesman 2004) to a broader class of input states and should be of independent interest.

Our theoretical method for proving these results is to construct a classical, local hidden variable model for the subtheory of quantum theory that consists of the stabilizer formalism and then determine the scope of additional quantum resources that are also described by this model. Indeed, our local hidden variable model is a distinguished quasi-probability representation with non-negative elements. For a d -dimensional quantum system there are many possible ways to represent arbitrary quantum states as quasi-probability distributions over a phase space of d^2 points and projective measurements as conditional quasi-probability distributions over the same space (see Ferrie and Emerson (2009) and Ferrie *et al* (2010) for further details). Perhaps unsurprisingly, it has been shown that the full quantum theory cannot be represented with non-negative elements in any such representation (Spekkens 2008, Ferrie and Emerson 2008, 2009, Ferrie *et al* 2010). However, one might expect that a subtheory of quantum theory that is

inadequate for quantum speed-up might be represented non-negatively, i.e. as a true probability theory, in some natural choice of quasi-probability representation. For the context described above, we seek a quasi-probability representation reflecting our natural operational restriction; in particular, we require that stabilizer states and projective measurements onto stabilizer states have a non-negative representation and that unitary stabilizer operations (i.e. Clifford transformations) correspond to stochastic processes. Conveniently, for quantum systems with odd Hilbert space dimension such a representation is already known to exist: this is the discrete Wigner function picked out by Gross (2006, 2007) from the broad class defined by Gibbons *et al* (2004). In such a representation it is natural to examine whether the resources that are necessary or sufficient for quantum speed-up correspond to those that are not represented by non-negative elements of the representation. Since the discrete Wigner function construction used here is only defined for systems of odd dimension, our results do not necessarily hold for qubits. We discuss the details of this peculiarity in the final section of this paper.

With this insight in hand the results of this paper may now be stated more carefully.

Classically efficient simulation of positive Wigner functions. The set of fault-tolerant quantum logic gates in the stabilizer formalism are known as the *Clifford* gates. Our first contribution is an explicit simulation protocol for quantum circuits composed of Clifford gates acting on input states with positive discrete Wigner representation. We also allow arbitrary product measurements with positive discrete Wigner representation. This simulation is efficient (linear) in the number of input registers to the quantum circuit. This simulation scheme is an extension of the celebrated Gottesman–Knill theorem and should be of independent interest.

Negativity is necessary for magic state distillation. This simulation protocol implies that states outside the stabilizer formalism with positive discrete Wigner function (bound universal states) are not useful for magic state distillation. The second contribution of this paper is to give a direct proof of this fact exploiting only the observation that negative discrete Wigner representation cannot be created by stabilizer operations. This proof has a more general range of applicability than the efficient simulation scheme and also makes clear the conceptual importance of negative quasi-probability as a *resource* for stabilizer computation.

Geometry of positive Wigner functions. The set of quantum states with positive discrete Wigner function strictly contains the set of (convex combinations of) stabilizer states. To prove this fact we determine the geometry of the region of quantum state space with positive discrete Wigner representation. Concretely, we show that the facets of the classical probability simplex defining the discrete Wigner function are also facets of the polytope with the (pure) stabilizer states as its vertices. Since there are many more facets of the stabilizer polytope than of the simplex, this suffices to show the existence of non-stabilizer states with positive representation.

This paper concludes with a discussion of this work and some avenues for future exploration.

2. Previous work

The Gottesman–Knill theorem provides an efficient classical simulation protocol for circuits of Clifford unitaries acting on stabilizer states. This result deals with pure qubit stabilizer state inputs and simulates the evolution of the full quantum state. The simulation scheme of the present paper deals with odd-dimensional systems, makes no distinction between mixed state and pure state input and allows the simulation of a large class of non-stabilizer states. However,

our scheme constructs a classical circuit with the same outcome probabilities as the quantum circuit and does not recover the evolution of the full quantum state.

A number of papers have addressed the question of which ancilla states enable universal quantum computation for the magic state model in qubit systems (Reichardt 2005, 2009a, 2009b, Campbell and Browne 2009, 2010, Campbell 2011, Virmani and Ratanje 2011, van Dam and Howard 2011). The most directly comparable result is the demonstration by Campbell and Browne (2010) that for any protocol on the input $\rho^{\otimes n}$ there exists a ρ outside the convex hull of stabilizer states that maps to a convex combination of stabilizers. As n grows these states are known to exist only within some arbitrarily small distance ϵ of the convex hull of stabilizer states. In contrast, the present result implies the existence of states a fixed distance from the hull which are not distillable by any protocol.

The present result is complementary to previous work connecting negativity in discrete Wigner function-type representations to quantum computational speed-up (Galvão 2005, Cormick *et al* 2006, van Dam and Howard 2011). In particular, van Dam and Howard (2009) have used techniques of this type to derive a bound on the amount of depolarizing noise a state can withstand before entering the stabilizer polytope. Their work deals only with prime-dimensional systems, and in this case it turns out that the noise threshold they derive is the same as the amount of noise required for their ‘maximally robust’ state to enter the region of positive states.

3. Background

3.1. The stabilizer formalism

Known schemes for fault-tolerant quantum computation allow for only a limited set of operations to be implemented directly on the encoded quantum information. For most known fault-tolerance schemes this restricted set is the stabilizer operations consisting of preparation and measurement in the computational basis and a restricted set of unitary operations. This restricted set is the Clifford group, and we now review the important parts of its structure for qudit systems (Gross *et al* 2007). The primitive object about which the stabilizer formalism is built is the Heisenberg–Weyl group, an extension of the qubit Pauli group to odd-dimensional systems. We will define the Heisenberg–Weyl group in terms of generalized X and Z operators. In odd *prime* dimension d these are given by their action on computational basis states,

$$\begin{aligned} X|x\rangle &= |x+1 \bmod d\rangle \\ Z|x\rangle &= \omega^x|x\rangle, \end{aligned}$$

where $\omega = \exp\left(\frac{2\pi i}{d}\right)$ is a d th primitive root of unity. The Heisenberg–Weyl operators are the d^2 operators generated by X and Z , which have a group structure if we include phases. A general element of the group is given as

$$T_{(a_1, a_2)} = \omega^{-\frac{a_1 a_2}{2}} Z^{a_1} X^{a_2} \in \mathbb{Z}_d \times \mathbb{Z}_d.$$

\mathbb{Z}_d is the finite field of d elements, and the field $\mathbb{Z}_d \times \mathbb{Z}_d$ will form the ‘phase space’ underpinning the discrete Wigner function defined below.

This definition applies only for prime dimension, but is easily promoted to arbitrary odd dimension. In this case the Heisenberg–Weyl operators are defined to be tensor products of the

Heisenberg–Weyl operators of the factor spaces. For a system with composite Hilbert space $H_a \otimes H_b \otimes \dots \otimes H_u$, the Heisenberg–Weyl operators may be written as

$$T_{(a_1, a_2) \oplus (b_1, b_2) \dots \oplus (u_1, u_2)} \equiv T_{(a_1, a_2)} \otimes T_{(b_1, b_2)} \otimes \dots \otimes T_{(u_1, u_2)}.$$

If the dimension of H_j is d_j the vector (j_1, j_2) is an element of $\mathbb{Z}_{d_j} \times \mathbb{Z}_{d_j}$ and the vector $(a_1, a_2) \oplus (b_1, b_2) \oplus \dots \oplus (u_1, u_2)$ is an element of $(\mathbb{Z}_{d_a} \times \mathbb{Z}_{d_a}) \times (\mathbb{Z}_{d_b} \times \mathbb{Z}_{d_b}) \times \dots \times (\mathbb{Z}_{d_u} \times \mathbb{Z}_{d_u})$. We denote the Heisenberg–Weyl group for dimension d as \mathcal{W}_d .

The Clifford operators, \mathcal{C}_d , are the set of unitary operators that map the Heisenberg–Weyl operators to the Heisenberg–Weyl operators under conjugation:

$$U \in \mathcal{C}_d \iff UT_u U^\dagger \in \mathcal{W}_d \forall T_u \in \mathcal{W}_d.$$

This is the normalizer of the Heisenberg–Weyl group. This group has many interesting features. For instance, operations of the Clifford group on computational basis input states can be efficiently simulated even though they create large amounts of entanglement.

For this paper it will be necessary to understand the Clifford group in terms of its representation over finite fields. Because it enormously simplifies the presentation we will restrict ourselves to working in the case when the system is composed of n subsystems with a common Hilbert space dimension p a prime. However, many of the important results carry over for arbitrary odd-dimensional systems (Gross 2006). The main result that we want is that a Clifford operation $U_{F,a} \in \mathcal{C}_{p^n}$ can be specified as,

$$U_{F,a} T_u U_{F,a}^\dagger = T_{Fu+a}.$$

That is, any Clifford unitary is uniquely specified by its action on the Pauli group and this is given by induced action on the label $\mathbf{u} \in (\mathbb{Z}_p \times \mathbb{Z}_p)^n$. The matrix F is a $2n \times 2n$ symplectic matrix with entries in \mathbb{Z}_p and $\mathbf{a} \in (\mathbb{Z}_p \times \mathbb{Z}_p)^n$. Clifford operations generally factor as

$$U_{F,a} = U_F T_a,$$

where $F \in \text{Sp}(2n, \mathbb{Z}_p)$, the group of symplectic matrices of size $2n$ with entries in \mathbb{Z}_p . A matrix is symplectic if it preserves the symplectic product, which may be defined in a natural way on finite fields. This structure is not important for the present paper so we do not cover it in detail. What is important is that this representation of the Clifford group has a size linear in n and thus can be easily tracked by a classical computer; this is at the heart of simulation results about the Clifford group.

Finally, we define stabilizer states to be any state that can be prepared by applying a Clifford unitary to a computational basis state. These states are important because they are the only pure states that can be prepared using our restricted operation set.

3.2. Magic state distillation

It is possible to implement stabilizer operations fault tolerantly, but these operations do not suffice for universal quantum computation. To promote stabilizer computation to universal quantum computation, some additional resource is required. This additional resource will be subject to large amounts of noise, so the question becomes: which non-stabilizer resources can be used to promote stabilizer computation and how can this be done? The first of these questions is the subject of this paper. The second question finds a particularly elegant solution in the form of magic state distillation (Bravyi and Kitaev 2004, Anwar *et al* 2012, Campbell *et al* 2012).

Magic state distillation protocols aim to consume a large number of copies of a non-stabilizer qudit input state ρ_{in} to produce a single non-stabilizer qudit output state ρ_{out} with higher fidelity to some non-stabilizer pure state. This output state is then consumed to implement some non-Clifford unitary gate. These protocols have the following structure:

- Prepare a number of copies of the input state $\rho_{\text{in}}^{\otimes n}$.
- Perform some Clifford operation on $\rho_{\text{in}}^{\otimes n}$.
- Measure the Heisenberg–Weyl observables on the last $n - 1$ registers and post-select on the outcome.

When these protocols succeed, the first register will be the output state ρ_{out} . Typically, these protocols work iteratively, repeatedly consuming $\rho_{\text{in}}^{\otimes n}$ until n copies of ρ_{out} have been produced and then using $\rho_{\text{out}}^{\otimes n}$ as the input to the protocol to produce ρ'_{out} and so on. The protocols we deal with here encompass but do not require this iterative structure.

It is not clear if for all input states ρ_{in} there exists a protocol to produce an output state with arbitrarily high fidelity to some non-stabilizer input state. We call states for which such a protocol exists ‘distillable’, and one of our results is to show that not all non-stabilizer states are distillable.

3.3. Quasi-probability representations

There is a long history of studying negativity in quasi-probability representations. The most notable examples come from quantum optics where the Wigner function (Kenfack and Życzkowski 2004) and the Q and P functions (Mandel 1986) play prominent roles. However, such approaches typically suffer in significance due to the problem of non-uniqueness of the choice of representation. While a quantum state may correspond to a negative valued quasi-probability function in one choice of quasi-probability representation, in another choice that same state can be positive and hence a valid classical probability density. In Ferrie *et al* (2010), two of us proved that for any choice of quasi-probability representation in which both quantum states and measurements are represented, at least some of the states and measurements must take on negative values. However, this result still leaves open the possibility that certain subsets of quantum states and measurements may be represented positively leading to a classical probability model for the corresponding subset of quantum operations (Wallman and Bartlett 2012). When the restricted subtheory is prescribed by an operational restriction, this question takes on a precise and relevant meaning. Indeed, such an approach has been considered already by Schack and Caves, who constructed classical probability models for few-qubit NMR experiments (Schack and Caves 1999) and thereby stimulated an important discussion of what kinds of resources might be required for universal quantum computation.

Our approach is to exploit the freedom in choice of quasi-probability representations (Ferrie 2011) in order to align the positive subtheory with the operational restriction defining the error-free resources in the magic state model. We seek a quasi-probability representation for which stabilizer states and projective measurements onto stabilizer states have positive representation and for which stabilizer transformations correspond to stochastic processes. This first step is easy given that such a representation already exists; this distinguished representation is the discrete Wigner function picked by Gross (2006, 2007) from the broad class defined by Gibbons *et al* (2004).

3.4. The discrete Wigner function

Our approach is to represent Clifford operations as stochastic processes over a discrete phase space. Intuitively, if the dynamics of the quantum system admit a representation as a classical statistical process, then it should not be sufficient for universal quantum computation. To that end, we look for a quasi-probability representation for quantum theory where stabilizer resources are represented positively. This is the discrete Wigner function.

The discrete Wigner representation of a state $\rho \in L(\mathbb{C}^{d^n})$ is a quasi-probability distribution over $\mathbb{Z}_d \times \mathbb{Z}_d$, which can be thought of as a $d \times d$ grid. This grid is the discrete analogue of the phase space of classical mechanics. The map taking quantum states to quasi-probability distributions on discrete phase space is uniquely specified by a set of *phase space point operators* $\{A_u\}$ (defined below). For each point \mathbf{u} in the discrete phase space there is a corresponding operator A_u and the value of discrete Wigner representation of ρ at this point is given as

$$W_\rho(\mathbf{u}) = \frac{1}{d} \text{Tr}(A_u \rho).$$

A quantum measurement with POVM $\{E_k\}$ is represented by assigning conditional (quasi-) probability functions over the phase space to each measurement outcome,

$$W_{E_k}(\mathbf{u}) = \text{Tr}(A_u E_k).$$

In the case when $W_{E_k}(\mathbf{u}) \geq 0 \forall \mathbf{u}$, this can be interpreted classically as the probability of getting outcome k given that the system is actually at point \mathbf{u} , $W_{E_k}(\mathbf{u}) = \text{Pr}(\text{outcome } k | \text{location } \mathbf{u})$. If both $W_\rho(\mathbf{u})$ and $W_{E_k}(\mathbf{u})$ are positive, then the law of total probability gives the probability of getting outcome k from a measurement of state ρ ,

$$\text{Pr}(k) = \sum_{\mathbf{u}} W_\rho(\mathbf{u}) W_{E_k}(\mathbf{u}).$$

In fact, this prediction reproduces the Born rule even when $W_\rho(\mathbf{u})$ or $W_{E_k}(\mathbf{u})$ takes on negative values.

We say a state ρ has positive representation if $W_\rho(\mathbf{u}) \geq 0 \forall \mathbf{u} \in \mathbb{Z}_d^n \times \mathbb{Z}_d^n$ and negative representation otherwise. We will say a measurement with POVM $M = \{E_k\}$ has positive representation if $W_{E_k}(\mathbf{u}) \geq 0 \forall \mathbf{u} \in \mathbb{Z}_d^n \times \mathbb{Z}_d^n, \forall E_k \in M$ and negative representation otherwise.

The phase space point operators are defined in terms of the Heisenberg–Weyl operators as

$$A_0 = \sum_{\mathbf{u}} T_{\mathbf{u}}, \quad A_{\mathbf{u}} = T_{\mathbf{u}} A_0 T_{\mathbf{u}}^\dagger.$$

These operators are Hermitian, so the discrete Wigner representation is real-valued. There are d^2 such operators for d -dimensional Hilbert space; they are informationally complete and orthogonal in the sense that $\text{Tr}(A_{\mathbf{u}} A_{\mathbf{v}}) = d \delta(\mathbf{u}, \mathbf{v})$.

These operators have several important features reflecting the salient properties of the discrete Wigner representation (Gibbons *et al* 2004, Gross 2006):

1. (Discrete Hudson’s theorem) If $|S\rangle$ is a stabilizer state then $\text{Tr}(A_{\mathbf{u}} |S\rangle\langle S|) \geq 0 \forall \mathbf{u}$, and the stabilizer states are the only pure states satisfying this property. That is, a pure state has positive representation if and only if it is a stabilizer state.

2. Clifford operators have the action $U_{F,a} A_u U_{F,a}^\dagger = A_{Fu+a}$. This means that

$$W_{U_{F,a}\rho U_{F,a}^\dagger}(v) = W_\rho(F^{-1}(v - a)),$$

so that Clifford transformations map to permutations of the underlying phase space and, in particular, Clifford operations preserve positive representation.

3. For $\rho = \sum_u p_u A_u$ and $\sigma = \sum_u q_u A_u$ the trace inner product is $\text{Tr}(\rho\sigma) = d \sum_u p_u q_u$.
4. The phase point operations in dimension d^n are tensor products of n copies of the d dimension phase space point operators.

4. Negative discrete Wigner representation is necessary for computational speed-up

We now establish that any quantum computation consisting of stabilizer operations acting on product input states with positive representation cannot produce an exponential computational speed-up. To this end, we give an explicit efficient classical simulation protocol for such circuits. Like the Gottesman–Knill protocol our scheme allows for the simulation of pure state stabilizer inputs to circuits composed of Clifford transformations and stabilizer measurements. However, our simulation scheme extends the Gottesman–Knill result in several ways. Firstly, it applies to systems of qudits rather than qubits. Secondly, it applies to mixed state inputs. Thirdly, and most remarkably, it applies to some non-stabilizer resources—namely those with positive discrete Wigner representation.

Any particular run of a quantum algorithm on n registers will produce a string \mathbf{k} of n measurement outcomes. These outcomes occur at random and we assign the random variable K_{quant} to be the algorithm output. The algorithm can then be considered as a way of sampling outcomes according to the distribution $\text{Pr}(K_{\text{quant}} = k)$. To simulate a quantum algorithm it suffices to give a simulating algorithm which samples from the distribution $\text{Pr}(K_{\text{quant}} = \mathbf{k})$, which is what we do here. Note that this form of simulation does not allow us to actually infer the distribution of outcomes, but it does suffice for many important tasks (for example, estimating the expected outcome).

The type of algorithms we treat here take the following form (see figure 1 for an example).

Algorithm Class 1. Family of simulable quantum algorithms

Algorithms in this class sample strings of measurement outcomes \mathbf{k} according to the distribution $\text{Pr}(K_{\text{quant}} = \mathbf{k})$ determined by the Born rule.

1. Prepare an initial n qudit input state $\rho_1 \otimes \dots \otimes \rho_n \in \rho \in L(\mathbb{C}^{p^n})$ where ρ_1, \dots, ρ_n have positive discrete Wigner representation.
 2. Until all registers have been measured:
 - (a) Apply a Clifford unitary gate U_F , labeled by the symplectic transformation $F \in \text{Sp}(2n, d)$.
 - (b) Measure the final qudit register using a measurement with positive discrete Wigner representation. Record the outcome k_j of the measurement of the j th register. Further steps in the computation may be conditioned on the outcome of this measurement.
-

Note that there is no loss of generality in considering only symplectic Clifford transformations, as the Heisenberg–Weyl component can be rolled into the measurement.

The essential idea of the simulation is to take seriously the hidden variable model the restrictions allow us. In the discrete Wigner picture, the system begins at point \mathbf{u} in the discrete phase space, which is unknown but definite and fixed. The effect of U_F is to move the system from the point \mathbf{u} to the point $F\mathbf{u}$, and measurement amounts to checking some region of the phase space to see if it contains the system. Since the vector \mathbf{u} and matrix F are size $2n$ with entries from \mathbb{Z}_d it is computationally efficient to classically store and update the system's location. Of course, a (positively represented) quantum state corresponds to a probability density over the space, so we must treat this a little more carefully. The simulation protocol is as follows.

Algorithm Class 2. Classical simulation algorithm

Algorithms in this class sample strings of measurement outcomes \mathbf{k} according to the distribution $\Pr(K_{\text{class}} = \mathbf{k})$

1. Sample $\mathbf{u} \in \mathbb{F}_d^{2n}$ according to the distribution $W_{\rho_1 \otimes \dots \otimes \rho_n}(\mathbf{u}) = W_{\rho_1}(\mathbf{u}_1)W_{\rho_2}(\mathbf{u}_2) \dots W_{\rho_n}(\mathbf{u}_n)$.
 2. Repeat until all registers have been measured:
 - (a) If the unitary U_F is applied, then update $\mathbf{u} \rightarrow F\mathbf{u}$.
 - (b) If the measurement M with corresponding POVM $\{E_k\}$ is made on the last register of the quantum circuit then report outcome k with probability $W_{E_k}(\mathbf{u}_m)$ where \mathbf{u}_m is the ontic position of the last qudit system, defined by $\mathbf{u} = \mathbf{u}_1 \oplus \mathbf{u}_2 \dots \oplus \mathbf{u}_m$. If the quantum algorithm conditions further steps on the outcome of measurement on this register then condition further steps of the simulation on measurement outcome k .
-

Our claim is that the classical algorithm in algorithm class 2 efficiently simulates the corresponding quantum algorithm in algorithm class 1. More precisely:

Theorem 1. *An n qudit quantum algorithm belonging to algorithm class 1 is simulatable by the corresponding $2n$ dit classical algorithm in algorithm class 2 in the sense that the distribution of outcomes \mathbf{k} is the same for both algorithms, $\Pr(K_{\text{class}} = \mathbf{k}) = \Pr(K_{\text{quant}} = \mathbf{k})$.*

Proof. The input to the classical circuit is a $2n$ dit string and the transformations are all matrices of size $2n$ with entries in \mathbb{Z}_d , so the $2n$ dit portion of the claim is obvious.

To show that this protocol genuinely simulates the circuit, it suffices to show that any string of measurement outcomes $\mathbf{k} = (k_1 k_2 \dots k_n)$ occurs with the same probability for both the original circuit and the simulation. Let us first consider probability distribution $\Pr(k_n)$ of the outcomes of the first measurement. In the quantum circuit, the preparation $\rho_1 \otimes \dots \otimes \rho_n$ is passed to the (possibly identity) gate U_F and the measurement M_n with the corresponding POVM $\{E_{k_n}\}$ is applied to the n th register. The probability of getting outcome k_n is then

$$\begin{aligned}
 \Pr_{\text{quant}}(k_n) &= \text{Tr}(U_F \rho_1 \otimes \dots \otimes \rho_n U_F^\dagger \mathbb{I} \otimes \dots \otimes \mathbb{I} \otimes E_{k_n}) \\
 &= \sum_{\mathbf{v} \in \mathbb{Z}_d^{2n}} W_{U_F \rho_1 \otimes \dots \otimes \rho_n U_F^\dagger}(\mathbf{v}) W_{\mathbb{I} \otimes \dots \otimes \mathbb{I} \otimes E_{k_n}}(\mathbf{v}) \\
 &= \sum_{\mathbf{v} \in \mathbb{Z}_d^{2n}} W_{\rho_1 \otimes \dots \otimes \rho_n}(F^{-1} \mathbf{v}) W_{\mathbb{I} \otimes \dots \otimes \mathbb{I} \otimes E_{k_n}}(\mathbf{v}), \tag{1}
 \end{aligned}$$




Figure 1. Example of a stabilizer circuit: ρ_i have positive representation, C_i are Clifford gates and M_i have positive representation. The choice of gate C_3 can be conditioned on the outcome of measurement M_3 .

where we have recast the inner product into the discrete Wigner form for convenience of comparison. We must now establish that the classical circuit has the same distribution.

Classically, if the system is initially at point \mathbf{v} on the discrete phase space, then probability of getting outcome k_n from the simulation circuit is given by

$$\begin{aligned} \Pr_{\text{class}}(k_n | \mathbf{v} \text{ sampled initially}) &= \Pr_{\text{class}}(k_n | \mathbf{F}\mathbf{v} \text{ final location}) \\ &= W_{\mathbb{I} \otimes \dots \otimes \mathbb{I} \otimes E_{k_n}}(\mathbf{F}\mathbf{v}), \end{aligned}$$

which just says that the system is moved from point \mathbf{v} to point $\mathbf{F}\mathbf{v}$ and the probability of outcome k_n is the probability that we see the system when we look at the region of phase space measured by E_{k_n} , which is $W_{E_{k_n}}(\mathbf{F}\mathbf{v})$ by definition. The total probability of outcome k_n is then

$$\begin{aligned} \Pr_{\text{class}}(k_n) &= \sum_{\mathbf{v} \in \mathbb{Z}_d^{2n}} \Pr_{\text{class}}(\mathbf{k} | \mathbf{v} \text{ sampled init.}) \Pr(\mathbf{v} \text{ sampled init.}) \\ &= \sum_{\mathbf{v} \in \mathbb{Z}_d^{2n}} W_{E_{k_1} \otimes \dots \otimes E_{k_n}}(\mathbf{F}\mathbf{v}) W_{\rho_1 \otimes \dots \otimes \rho_n}(\mathbf{v}) \\ &= \sum_{\mathbf{v} \in \mathbb{Z}_d^{2n}} W_{E_{k_1} \otimes \dots \otimes E_{k_n}}(\mathbf{v}) W_{\rho_1 \otimes \dots \otimes \rho_n}(\mathbf{F}^{-1}\mathbf{v}). \end{aligned} \quad (2)$$

Comparing algorithm class 1, the distribution of measurement outcomes on the last register for the quantum circuit, and algorithm class 2, the simulated distribution of measurement outcomes on the last register, we see that they are the same.

If the quantum algorithm is independent of the measurement outcomes, then simply applying the above argument to each register would suffice to complete the proof. However, in general, adaptive schemes are possible, such as the algorithm illustrated in figure 1 where the final gate applied depends on the outcome of the measurement on the third qudit. Using the assumption that the registers are measured from the last to the first, we can factor the distribution of outcome strings as

$$\Pr(\mathbf{k}) = \Pr(k_1 | k_2 \dots k_n) \Pr(k_2 | k_3 \dots k_n) \dots \Pr(k_{n-1} | k_n) \Pr(k_n).$$

Since the simulation conditions on measurement outcome in exactly the same way as the original quantum algorithm, a simple inductive argument shows that the distribution of outcomes must be the same for the quantum algorithm and its classical simulator. \square

Corollary 2. *Quantum algorithms belonging to algorithm class 1 offer no super-linear advantage over classical computation.*

Proof. We have seen that if it is computationally efficient (linear in the number of qudits) to sample from the classical distributions corresponding to the input state and the measurements, then such quantum circuits are efficiently simulable. Since we have assumed separability of the input and measurements and the discrete Wigner function factors, this efficient sampling is guaranteed. \square

A couple of remarks are in order. We have restricted ourselves to separable inputs and measurements, but this is not strictly necessary for efficient simulation. Any positively represented preparation or measurement can be accommodated provided it is possible to classically efficiently sample from the corresponding distribution. Since it is exponentially difficult to even write down general quantum states, this is a rather strong restriction.

Also, note that our simulation protocol only samples from the output distribution of a circuit, whereas the Gottesman–Knill protocol gives the full quantum state output in the case when the input is a pure stabilizer state. It appears that the present protocol is weaker in this respect. However, the discrete Wigner functions of pure stabilizer states are uniformly valued lines on the phase space (Gross 2006) and these are fully specified by only two points. If we are sure that the input state to a Clifford circuit is a stabilizer state, then we can sample two distinct points from the corresponding distribution and determine where the circuit maps them. These two output points then suffice to fix the line corresponding to the output stabilizer state.

Finally, we note that, in the context of magic state distillation for example, one may increase the size of the input register conditional on measurement outcomes. This can be accounted for in the simulation protocol above by simply increasing the size of the phase space accordingly and sampling from the new additional positive Wigner functions.

5. Magic state distillation

The main significance of the simulation result just established is that, for noisy preparation and measurement, it is possible to extend the efficient simulation of quantum circuits beyond the purview of the stabilizer formalism. This result is of major theoretical importance, but it also has practical significance. In particular, the simulation scheme addresses the magic state model, which supplements error-free stabilizer resources with high-fidelity additional gates produced through the consumption of non-stabilizer ancillas. Recall that the backbone of this process is a distillation protocol that uses stabilizer resources applied to a large number of ancilla input states to produce a few highly pure non-stabilizer states. An immediate corollary of our simulation protocol is that states with positive discrete Wigner representation are not useful for computational speed-up in the magic state model. As we will make explicit in the next section (see also figures 2–4), this includes a large class of states outside the stabilizer formalism. Thus, we have resolved the long-standing open problem of whether *all* non-stabilizer states promote stabilizer computation to universal quantum computation (Bravyi and Kitaev 2005).

The class of algorithms in the previous section encompasses a large variety of magic state distillation protocols, but it is still conceptually unclear why states with positive discrete Wigner representation are not useful for magic state distillation. This is especially true since we typically think of the outcome of a magic state distillation routine as a quantum state, rather than a string of measurement outcomes as in the simulation protocol. If we did keep track of the full input distribution, then we would be able to use it to reconstruct the quantum state output of a distillation procedure; unfortunately, this is impossible to do efficiently, but a little introspection




Figure 2. A cartoon of the intersection of the discrete Wigner probability simplex (the triangular region) with the quantum state space (the circle). The simplex intersects the boundary at stabilizer states (bold dots). The region of convex combinations of stabilizer states is strictly contained within the set of quantum states that also lie inside the simplex. The quantum states outside the simplex are the bound states. Finally, the quantum states with negative discrete Wigner representation are those lying outside the positive discrete Wigner simplex. We show that the half-space inequalities defining the facets of discrete Wigner simplex also define the facets of the stabilizer polytope, a fact reflected in this cartoon.

shows it is actually not necessary. We do not need to know the final quantum state; we only need to know that it is not helpful in doing quantum computation. The simulation protocol makes it clear that if the quantum state that is put into the circuit has a Wigner function that is a genuine probability distribution, then the quantum state that is output (which we measure) must also correspond to a genuine probability distribution. Inspired by this observation and in view of the great importance of magic state protocols, we devote this section to a direct proof that negative discrete Wigner representation of the ancilla resource states is necessary for such states to be distillable (using stabilizer resources) to a non-stabilizer state of arbitrary purity.

The essential insight of the proof used here is that negative discrete Wigner representation is a resource that cannot be created using stabilizer operations; if the input states to a distillation protocol have no negativity in its discrete Wigner representation, then the output will not either. This resource character is one of the major insights of this work. Beyond its conceptual value the alternative proof presented below also closes several loopholes and alternative models not addressed by the simulation protocol, which we discuss at the end of the section.

Conventional magic state protocols perform a Clifford unitary on the input state $\rho_{\text{in}}^{\otimes n} \in L(\mathbb{C}_d \otimes \mathbb{C}_d^{n-1})$ and make a computational basis measurement on the final $n - 1$ qudits, post selecting on the $|0\rangle$ outcome. This outputs the state

$$\rho_{\text{out}} = \frac{\text{Tr}_{\mathbb{C}_d^{n-1}} (\mathbb{I} \otimes |0\rangle\langle 0|^{\otimes n-1} U \rho_{\text{in}}^{\otimes n} U^\dagger)}{\text{Normalization}},$$

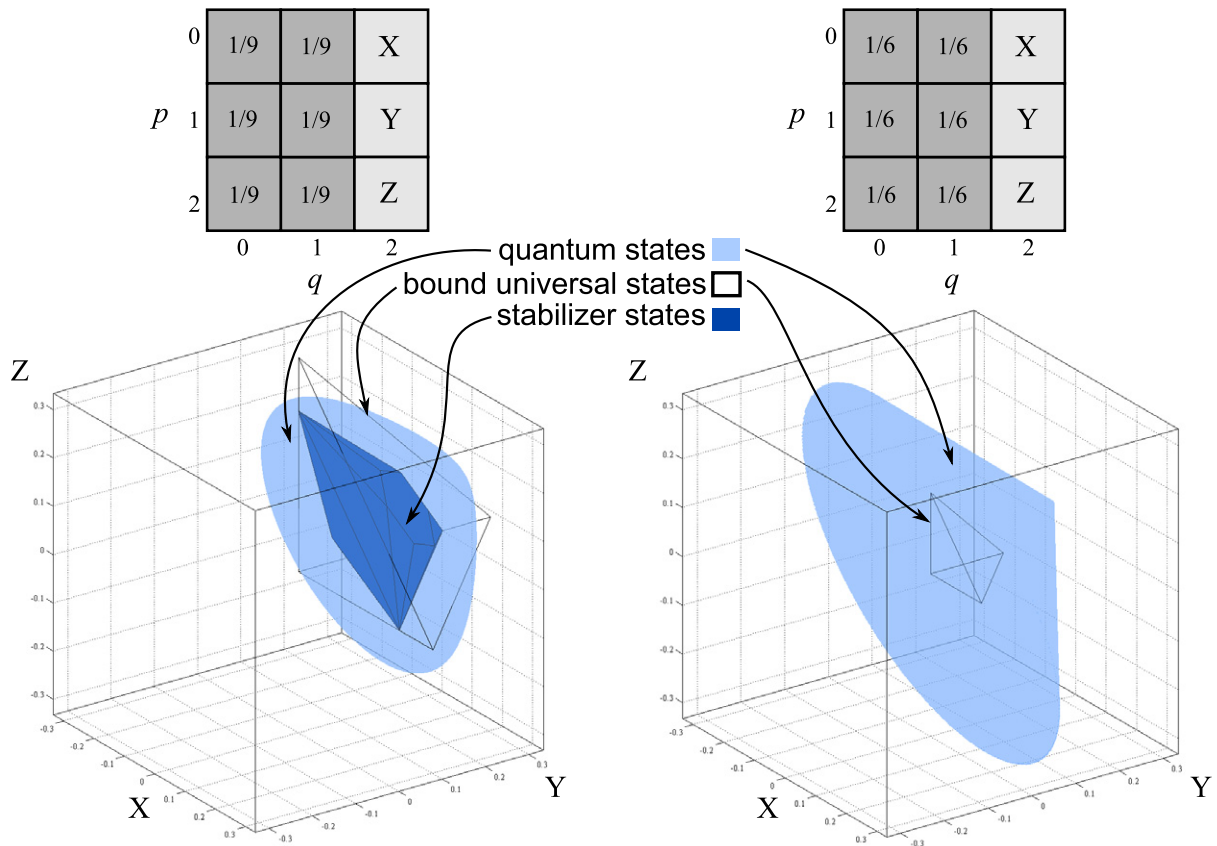


Figure 3. Orthogonal 3D slices of qutrit state space. Above each slice are the six values of the Wigner function which are fixed at a value of $1/9$ (left) and $1/6$ (right). The remaining three values are allowed to vary and carve out regions depicted in the graphs. In each case, due to symmetries, there are $\binom{9}{6} = 84$ such slices which are identical (up to a relabeling of the axes). Note that the slice on the right does not cut through the stabilizer polytope but does contain a region of bound states. Also note that this slice contains one of the nine states with a maximal negativity of $1/3$ while the slice on the left, and those equivalent up to permutations, are the only ones which feature the maximally mixed state $(X, Y, Z) = 1/9$. See also figure 4 for a 2D slice of the figure on the left.

where the partial trace knocks out the ancilla systems and the normalization in the denominator just guarantees that $\text{Tr}(\rho_{\text{out}}) = 1$. We examine significantly more general protocols: instead of requiring n copies of a qudit input state, we allow an arbitrary positively represented input state; in place of a Clifford unitary, we allow any completely positive map which preserves the set of positively represented states; and in place of the computational basis measurement, we allow any positively represented projective measurement. Since positive representation is convex, this suffices to eliminate classical randomness as a potential loophole. It also precludes choices of entangled stabilizer measurements, as these are still positively represented. For convenience of presentation, we define

$$F(\rho) = \min_u \text{Tr}(A_u \rho),$$




Figure 4. Orthogonal 2D slice of qutrit state space. On the left are the seven values of the Wigner function which are fixed at a value of $1/9$. The remaining two are allowed to vary. The maximally mixed state is the point $(X, Y) = (1, 1)/9$. The various regions carved out by varying these values are shown on the right. There are $\binom{9}{7} = 36$ such slices which are identical (up to a relabeling of the axes). These would be the only slices featuring the maximally mixed state. Note the similarity of the caricature in figure 2, remarkable since this cartoon was merely the intersection of the simplest simplex (a triangle) with the simplest continuous state space (a circle).

which is the maximally negative point of the quasi-probability representation of ρ over the phase space. If the input state to a distillation routine is positively represented (i.e. $F(\rho) \geq 0$), then its output is also positively represented.

Theorem 3. Let ρ_{in} be a density operator on an n qudit Hilbert space such that $F(\rho_{in}) \geq 0$. Let Λ be a (completely positive) map on this space for which $F(\rho) \geq 0 \Rightarrow F(\Lambda(\rho)) \geq 0$. Let P be a positively represented projector on this space. If ρ_{out} is produced by acting on ρ_{in} with Λ and post selecting a measurement on the last $n - 1$ qudits on outcome P , then $F(\rho_{out}) \geq 0$.

Proof. Since we can use Heisenberg–Weyl operations to cycle between phase point operations, without loss of generality

$$\begin{aligned}
 F(\tilde{\rho}) &= \min_u \text{Tr}(A_u \rho) \\
 &= \frac{\text{Tr}(A_0 \otimes \mathbb{I}^{\otimes n-1} \cdot \mathbb{I} \otimes P \cdot \Lambda(\rho_{in}) \cdot \mathbb{I} \otimes P)}{\text{Normalization}} \\
 &= \frac{\text{Tr}(A_0 \otimes P \cdot \Lambda(\rho_{in}))}{\text{Normalization}}.
 \end{aligned}$$

The denominator is always positive, so

$$F(\tilde{\rho}) \geq 0 \iff \text{Tr}(A_0 \otimes P \cdot \Lambda(\rho_{in})) \geq 0.$$

By assumption $F(\rho_{\text{in}})$ is positive and thus so is $F(\Lambda(\rho_{\text{in}}))$. We write $P = \sum_v z_v A_v$ where, since P is positively represented, $z_u \geq 0$ by the discrete Hudson's theorem. This gives us that

$$\text{Tr}(A_0 \otimes P \cdot \Lambda(\rho_{\text{in}})) = \sum_v z_v \text{Tr}(A_0 \otimes A_v \Lambda(\rho_{\text{in}})).$$

The non-negativity of $F(\Lambda(\rho_{\text{in}}))$ implies, by definition, that $\text{Tr}(A_0 \otimes A_v \Lambda(\rho_{\text{in}})) \geq 0$ so it must be the case that $\text{Tr}(A_0 \otimes P \cdot \Lambda(\rho_{\text{in}})) \geq 0$ and this implies that $F(\tilde{\rho}) \geq 0$. \square

This proof has a few technical merits over the simulation result of the previous section. Namely, it does not require the input state to the distillation protocol to be separable or otherwise efficiently sampleable and the update dynamics are not limited to only Clifford unitary maps. Indeed it is not even necessary to have a description of the input state or the transformation; all that is required is the promise that $\Lambda(\rho_{\text{in}})$ has positive discrete Wigner representation. Since many input states and channels do not admit any efficient representation, this means that there are resources that are provably useless for distillation even when an efficient simulation of the corresponding distillation process would be impossible.

6. The geometry of positively represented states

The Gottesman–Knill theorem already establishes that the action of (qubit) Clifford unitary operations on stabilizer states is efficiently classically simulatable. Since the only pure states with positive discrete Wigner representation are stabilizer states, it is natural to wonder if every positively represented state is a mixture of stabilizer states. As we have already alluded to, remarkably this is *false*. To establish this, we will clarify the geometry of the region of state space which has positive representation and show that it strictly contains the set of mixtures of stabilizer states. Combined with the results of the previous sections, this establishes our simulation protocol as an extension of Gottesman–Knill and proves the existence of bound states for magic state distillation, states which are not convex combinations of stabilizer states but which are nevertheless not distillable using perfect Clifford operations.

The set of convex combinations of stabilizer states is a convex polytope with the stabilizer states as vertices. Any polytope can be defined either in terms of its vertices or as a list of half-space inequalities called facets. Intuitively, these correspond to the faces of three-dimensional (3D) polyhedrons. We show that in powers of prime dimension each of the d^2 phase space, point operators define a facet of the stabilizer polytope. These are only a proper subset of the faces of the stabilizer polytope, implying the existence of states with positive representation which are not convex combinations of stabilizer states. See figure 2 for a cartoon capturing the intuition for this result.

The stabilizer polytope may be thought of as a bounded convex polytope living in \mathbb{R}^{d^2-1} , the space of d -dimensional mixed quantum states. A minimal half-space description for a polytope in \mathbb{R}^D is a finite set of bounding equalities called facets $\{F_i, f_i\}$ with $F_i \in \mathbb{R}^D$ and $f_i \in \mathbb{R}$. $X \in \mathbb{R}^D$ is in the polytope if and only if $X \cdot F_i \leq f_i \forall i$. In the usual quantum state space the vectors X of interest are density matrices, the inner product is the trace inner product and facets may be defined as $\{\hat{A}_i, a_i\}$ where \hat{A}_i are Hermitian matrices and

$$\rho \in \text{polytope} \iff \text{Tr}(\rho \hat{A}_i) \leq a_i, \forall i.$$

The objective is to show that $\{-A_u, 0\}$ are facets of the polytope defined by stabilizer state vertices.

It is possible to explicitly compute a facet description for a polytope given the vertex description, but the complexity of this computation scales polynomially as the number of vertices. Since the number of stabilizer states grows super-exponentially with the number of qudits (Gross 2006) the conversion is generally impractical. The analytic proof given here circumvents this issue. We also remark that the work of Cormick *et al* (2006) implies that the phase space point operators considered here are facets for the case of *prime* dimension.

Theorem 4. *The d^2 phase space point operators $\{A_u\}$ with the inequalities $\text{Tr}(\rho A_u) > 0$ define facets of the stabilizer polytope.*

Proof. To establish that a half-space inequality for a polytope in \mathbb{R}^D is a facet, there are two requirements: every vertex must satisfy the inequality and there must be a set of vertices saturating the inequality which span a space of dimension D (Ziegler 1995).

The requirement that all vertices satisfy the half-space inequality is $\text{Tr}(A_u S) \geq 0$ for every stabilizer state S , and this is the discrete Hudson's theorem.

We consider the stabilizer polytope as an object in \mathbb{R}^{d^2-1} and look for a set of $d^2 - 1$ linearly independent vertices which satisfy $\text{Tr}(A_u S) = 0$. Since we are restricting ourselves to power of prime dimension we may choose a complete set of mutually unbiased bases of $d(d+1)$ states from the full set of stabilizer states. Suppose that more than $d+1$ states V_i from this set satisfy $\text{Tr}(V_i A_u) > 0$. Then a counting argument shows that there must be two distinct states V_0, V_1 belonging to an orthonormal basis which satisfies this criterion. But then

$$\begin{aligned} \text{Tr}(V_0 V_1) &= \frac{1}{d} \sum_v \text{Tr}(V_0 A_v) \text{Tr}(V_1 A_v) \\ &\geq \frac{1}{d} \text{Tr}(V_0 A_u) \text{Tr}(V_1 A_u) \neq 0, \end{aligned}$$

which contradicts the orthonormality. Thus at least $d(d+1) - (d+1)$ states in the mutually unbiased bases satisfy $\text{Tr}(A_u V_i) = 0$. These are the required set of $d^2 - 1$ linearly independent vertices. \square

The phase space point operators considered here give only a proper subset of the defining half-space inequalities for the stabilizer polytope. This means that there are states that may not be written as a convex combination of stabilizer states which nevertheless satisfy $\text{Tr}(A_u \rho) \geq 0$ for all phase space point operators. That is, there are positive states which are not in the convex hull of stabilizer states. These are bound states for magic state distillation. An explicit example of such a state for the qutrit is given by Gross (2006). These regions can be visualized by taking 2D and 3D slices of the qutrit state space. Such slices are depicted in figures 3 and 4.

7. Discussion and conclusions

We have shown that for systems of power of odd prime dimension, a necessary condition for computational speed-up using Clifford unitaries is negativity of the discrete Wigner representation of the inputs. This result is immediately relevant in the context of magic state distillation, where it shows that a necessary condition for distillability is negative representation of the ancilla preparation. We have also shown that the phase space point operators defining the discrete Wigner function correspond to a privileged set of facets of the stabilizer polytope. Taken

together, the two results imply the existence of non-stabilizer resources which do not promote Clifford computation to universal quantum computation; and in particular, this establishes the existence of bound states for magic state distillation, or bound universal states.

We motivated the development of negative discrete Wigner representation by analogy to entanglement theory, with Clifford operations playing the role of local operations and classical communication and stabilizer states playing the role of separable states. It is known that there are slightly entangled mixed states that cannot be consumed by distillation routines to produce highly entangled states (Horodecki *et al* 1998). The non-stabilizer but positively represented quantum states are exactly analogous to these bound entangled states. Similarly, it is known that for pure states large amounts of entanglement are required for quantum computational speed-up (Vidal 2003), but for mixed states this is still an open question. However, for negative discrete Wigner representation there is no relevant distinction between mixed states and pure states. Moreover, although it is not yet known whether every negatively represented state is distillable we conjecture this to be the case.

The discrete Wigner function considered in this work is defined by analogy with the more familiar continuous variable Wigner function. It is natural to wonder if the results of this work extend to the infinite-dimensional case. In the continuous case a pure state has positive Wigner representation if and only if it is a Gaussian state (Hudson 1974) (like stabilizer states in the discrete case) and a unitary evolution acts as a symplectic flow on the phase space if and only if it corresponds to a quadratic Hamiltonian (Weedbrook *et al* 2012) (like Clifford gates in the discrete case). A result of Bröcker and Werner (1995) shows that there are mixed states that cannot be represented as probabilistic combinations of Gaussian states but that nevertheless have positive Wigner function. The question that naturally arises is then: is it possible to efficiently simulate quantum systems with quadratic Hamiltonian dynamics acting on non-Gaussian mixed states if these states have positive representation? The answer to this question, *yes*, is obtained by giving an explicit efficient simulation protocol for such systems (Veitch *et al* in preparation). This establishes that two of the main results of this paper (efficient simulation and non-stabilizer mixed states with positive discrete Wigner representation) extend to the continuous case. It is interesting to ask whether our third result, namely that negative discrete Wigner representation is a necessary resource for distillation, also has an analogue. That is, does the continuous variable case admit something analogous to the magic state model which allows noisy negatively represented states to be consumed in order to promote linear optics to full universal quantum computational power?

Of course, there remains a final detail that we have not yet addressed. The discrete Wigner function underpinning our analysis is only defined for odd-dimensional systems. Is it possible to find a similar construction for qubits? The discrete Wigner function used here has two crucial properties: Clifford operations are stochastic transformations on the underlying phase space and this phase space is separable. For qubit systems there is no known analogue. Indeed, it is easy to see that a quasi-probability representation defined by any proper subset of the facets of the qubit stabilizer polytope in a fashion analogous to what has been done here will assign positive representation to some subset of the magic states. The construction of the discrete Wigner function, and the Clifford group, relies critically on the mathematics of finite fields, and it is well known that fields of characteristic 2 behave fundamentally differently from fields of any other characteristic. This fact is reflected in the theory of error correction where somewhat different protocols are required for dealing with bits and qubits than those used for dits and qudits. In the case of error correction, although qubits require more involved mathematics,

the conceptual underpinnings are the same irrespective of the underlying dimensionality of the dits. It is not unreasonable to hope that a similar result will hold for qubits and that an appropriately modified mathematical strategy will preserve the conceptual insights and related technical results obtained in the qudit case. If this turns out not to be the case, then an exact understanding of why the model fails for qubits will undoubtedly provide deep insights into the workings of quantum theory and quantum computation.

The most interesting outstanding question raised by this work is whether the ability to prepare any state with negative discrete Wigner representation is sufficient for promoting Clifford computation to universal quantum computation. In *prime* dimension, the discrete Wigner construction is the *unique* choice of quasi-probability representation covariant under the action of Clifford operations (Gross 2006), where the law of total probability is required to hold. On this basis, we conjecture that the condition here is sufficient. From the work of Campbell *et al* (2012), it is already known that access to any non-stabilizer *pure* state (or equivalently any negatively represented pure state) suffices. If this conjecture is true, then this implies an equivalence of two previously unrelated concepts of non-classicality, namely quantum computational speed-up and negative quasi-probability representation.

Acknowledgments

We thank Nathan Weibe, Daniel Gottesman and Ben Reichardt for helpful comments. After the completion of this work we became aware of a complementary work by Eisert and Mari (2012) which develops an alternative efficient simulation for systems with positively represented Wigner function in odd and infinite dimensions. We acknowledge financial support from CIFAR, the Government of Canada through NSERC and the Ontario Government through OGS and ERA. JE thanks the Perimeter Institute for Theoretical Physics where this work was completed.

References

- Aaronson S and Gottesman D 2004 Improved simulation of stabilizer circuits *Phys. Rev. A* **70** 052328
- Anwar H, Campbell E T and Browne D E 2012 Qutrit magic state distillation *New J. Phys.* **14** 063006
- Bravyi S and Kitaev A 2004 Universal quantum computation with ideal Clifford gates and noisy ancillas *Phys. Rev. A* **71** 022316
- Bravyi S and Kitaev A 2005 Universal quantum computation with ideal Clifford gates and noisy ancillas *Phys. Rev. A* **71** 022316
- Bröcker T and Werner R F 1995 Mixed states with positive Wigner functions *J. Math. Phys.* **36** 62–75
- Campbell E T 2011 Catalysis and activation of magic states in fault-tolerant architectures *Phys. Rev. A* **83** 032317
- Campbell E and Browne D 2009 On the structure of protocols for magic state distillation *Theory of Quantum Computation, Communication and Cryptography (Lecture Notes in Computer Science vol 5906)* ed A Childs and M Mosca (Berlin: Springer) pp 20–32
- Campbell E T and Browne D E 2010 Bound states for magic state distillation in fault-tolerant quantum computation *Phys. Rev. Lett.* **104** 030503
- Campbell E T, Anwar H and Browne D E 2012 Magic state distillation in all prime dimensions using quantum Reed–Muller codes arXiv:1205.3104
- Chen X, Chung H, Cross A W, Zeng B and Chuang I L 2008 Subsystem stabilizer codes cannot have a universal set of transversal gates for even one encoded qudit *Phys. Rev. A* **78** 012353

- Cormick C, Galvao E F, Gottesman D, Paz J P and Pittenger A O 2006 Classicality in discrete Wigner functions *Phys. Rev. A: At. Mol. Opt. Phys.* **73** 012301
- Datta A, Flammia S T and Caves C M 2005 Entanglement and the power of one qubit *Phys. Rev. A* **72** 042316
- Douçot B and Vidal J 2002 Pairing of Cooper pairs in a fully frustrated Josephson-junction chain *Phys. Rev. Lett.* **88** 227005
- Eastin B and Knill E 2009 Restrictions on transversal encoded quantum gate sets *Phys. Rev. Lett.* **102** 110502
- Eisert J and Mari A 2012 Positive Wigner functions render classical simulation of quantum computation efficient arXiv:1208.3660
- Ferrie C 2011 Quasi-probability representations of quantum theory with applications to quantum information science *Rep. Prog. Phys.* **74** 116001
- Ferrie C and Emerson J 2008 Frame representations of quantum mechanics and the necessity of negativity in quasi-probability representations *J. Phys. A: Math. Theor.* **41** 352001
- Ferrie C and Emerson J 2009 Framed Hilbert space: hanging the quasi-probability pictures of quantum theory *New J. Phys.* **11** 063040
- Ferrie C, Morris R and Emerson J 2010 Necessity of negativity in quantum theory *Phys. Rev. A* **82** 044103
- Galvão E F 2005 Discrete Wigner functions and quantum computational speedup *Phys. Rev. A* **71** 042302
- Gibbons K S, Hoffman M J and Wootters W K 2004 Discrete phase space based on finite fields *Phys. Rev. A* **70** 062101
- Gottesman D 1997 Stabilizer codes and quantum error correction *PhD Thesis* California Institute of Technology (arXiv:quant-ph/9705052v1)
- Gottesman D 2006 *Quantum Error Correction and Fault-Tolerance* (Amsterdam: Elsevier) pp 196–201 (arXiv:quant-ph/0507174)
- Gross D 2006 Hudson's theorem for finite-dimensional quantum systems *J. Math. Phys.* **47** 122107
- Gross D 2007 Non-negative Wigner functions in prime dimensions *Appl. Phys. B* **86** 367–70
- Gross D, Audenaert K and Eisert J 2007 Evenly distributed unitaries: on the structure of unitary designs *J. Math. Phys.* **48** 052104
- Horodecki M, Horodecki P and Horodecki R 1998 Mixed-state entanglement and distillation: is there a 'bound' entanglement in nature? *Phys. Rev. Lett.* **80** 5239–42
- Hudson R 1974 When is the Wigner quasi-probability density non-negative? *Rep. Math. Phys.* **6** 249–52
- Kenfack A and Życzkowski K 2004 Negativity of the Wigner function as an indicator of non-classicality *J. Opt. B: Quantum Semiclass. Opt.* **6** 396
- Knill E and Laflamme R 1998 Power of one bit of quantum information *Phys. Rev. Lett.* **81** 5672–5
- Lloyd S 2002 Quantum computation with Abelian anyons *Quantum Inform. Process.* **1** 13–8
- Mandel L 1986 Non-classical states of the electromagnetic field *Phys. Scr.* **T12** 34
- Moore G and Read N 1991 Nonabelions in the fractional quantum hall effect *Nucl. Phys. B* **360** 362–96
- Reichardt B W 2005 Quantum universality from magic states distillation applied to CSS codes *Quantum Inform. Process.* **4** 251–64
- Reichardt B W 2009a Error-detection-based quantum fault-tolerance threshold *Algorithmica* **55** 517–56
- Reichardt B W 2009b Quantum universality by state distillation *Quantum Inform. Comput.* **9** 1030–52
- Schack R and Caves C M 1999 Classical model for bulk-ensemble NMR quantum computation *Phys. Rev. A* **60** 4354–62
- Spekkens R W 2008 Negativity and contextuality are equivalent notions of nonclassicality *Phys. Rev. Lett.* **101** 020401
- van Dam W and Howard M 2009 Tight noise thresholds for quantum computation with perfect stabilizer operations *Phys. Rev. Lett.* **103** 170504
- van Dam W and Howard M 2011 Noise thresholds for higher-dimensional systems using the discrete Wigner function *Phys. Rev. A* **83** 032310
- Veitch V, Wiebe N, Ferrie C and Emerson J 2012 Efficient simulation scheme for a class of quantum optics experiments with non-negative Wigner representation arXiv:1210.1783

- Vidal G 2003 Efficient classical simulation of slightly entangled quantum computations *Phys. Rev. Lett.* **91** 147902
- Virmani S and Ratanje N 2011 Generalised state spaces and non-locality in fault tolerant quantum computing schemes arXiv:1007.3455v2
- Wallman J and Bartlett S 2012 Non-negative subtheories and quasiprobability representations of qubits *Phys. Rev. A* **85** 062121
- Weedbrook C, Pirandola S, Garcia-Patron R, Cerf N J, Ralph T C, Shapiro J H and Lloyd S 2012 Gaussian quantum information *Rev. Mod. Phys.* **84** 621
- Zeng B, Cross A and Chuang I L 2007 Transversality versus universality for additive quantum codes arXiv:0706.1382
- Ziegler G M 1995 *Lectures on Polytopes* (Berlin: Springer)