



CHORUS

This is the accepted manuscript made available via CHORUS. The article has been published as:

Estimating Outcome Probabilities of Quantum Circuits Using Quasiprobabilities

Hakop Pashayan, Joel J. Wallman, and Stephen D. Bartlett
Phys. Rev. Lett. **115**, 070501 — Published 10 August 2015

DOI: [10.1103/PhysRevLett.115.070501](https://doi.org/10.1103/PhysRevLett.115.070501)

Estimating outcome probabilities of quantum circuits using quasiprobabilities

Hakop Pashayan,¹ Joel J. Wallman,² and Stephen D. Bartlett¹

¹*Centre for Engineered Quantum Systems, School of Physics,
The University of Sydney, Sydney, NSW 2006, Australia*

²*Institute for Quantum Computing and Department of Applied Mathematics,
University of Waterloo, Waterloo, Ontario, Canada, N2L 3G1*

(Dated: 12 July 2015)

We present a method for estimating the probabilities of outcomes of a quantum circuit using Monte Carlo sampling techniques applied to a quasiprobability representation. Our estimate converges to the true quantum probability at a rate determined by the total negativity in the circuit, using a measure of negativity based on the 1-norm of the quasiprobability. If the negativity grows at most polynomially in the size of the circuit, our estimator converges efficiently. These results highlight the role of negativity as a measure of non-classical resources in quantum computation.

Estimating the probability of a measurement outcome in a quantum process using only classical methods is a longstanding problem that remains of acute interest today. Directly calculating such probabilities using the Born rule is inherently inefficient in the size of the quantum system, and efficiently estimating such probabilities for a generic quantum process is expected to be out of reach of classical computers.

Nonetheless, there are interesting and nontrivial classes of quantum circuits for which we *can* efficiently estimate the probabilities of outcomes. The canonical example of such a class is that of stabilizer circuits. Such circuits can create highly-entangled states and perform many of the fundamental operations involved in quantum computing (teleportation, quantum error correction, distillation of magic states) but the celebrated Gottesman-Knill theorem allows such circuits to be classically simulated efficiently [1]. Other examples include fermionic linear optics/matchgates [2, 3], and some classes of quantum optics [4–6]. While these methods may be extended to include bounded numbers of operations outside of the class (for example, Ref. [1]), such extensions generally treat all operations outside of the class on an equal footing (for example, the cost of adding noisy magic states is the same as adding pure magic states) and so do not provide any insight into the relative resources of different operations.

In this Letter, we present a general method for estimating outcome probabilities for quantum circuits using quasiprobability representations. Simulation methods based on quasiprobability representations have a long history in physics [7], and have recently been used in quantum computation to identify classes of operations that are efficiently simulatable [8–10]. Our method allows for estimation in circuits wherein the quasiprobabilities may go negative. That is, while making the most efficient use of circuit elements that are represented non-negatively, it nonetheless provides an unbiased estimator of the true quantum outcome probability regardless of the inclusion of more general elements that are negatively represented. We quantify the performance of this

method by providing an upper bound on the rate of convergence of this estimator that scales with a measure of the total amount of negativity in the circuit.

Probability estimation.—Consider quantum circuits of the following form. The circuit initiates with N qudits (d -level quantum systems) in a product state, evolves through a circuit consisting of $L = \text{poly}(N)$ elementary gates that act nontrivially on at most a fixed number of qudits (for example, 1- and 2-qudit gates), and terminates with a product measurement, i.e., an independent measurement of each qudit. Universal quantum computation can be achieved with circuits of this form. Note that we do not include circuits with intermediate measurements and conditional operations based on their outputs (we return to this consideration in the discussion).

We aim to estimate the probability of a fixed outcome $\vec{o} = (o_1, \dots, o_N)$ where o_j denotes the outcome of the measurement on the j th qudit. (Note that *estimation* of the probability of a fixed outcome is distinct from a *simulation* as in Refs. [8, 9], wherein different outcomes are sampled from this distribution.) A natural benchmark for the precision of an estimator is the precision that can be obtained from sampling the quantum circuit itself. If we had access to a quantum computer that implemented a circuit in this class, then we could use it to estimate the probability of a fixed outcome by computing the observed frequency $f_s(\vec{o})$ of outcome \vec{o} over s samples. By the Hoeffding inequality, $f_s(\vec{o})$ will be within ϵ of the quantum probability $p(\vec{o})$ with probability $1 - \delta$ provided the number of samples $s(\epsilon, \delta)$ satisfies

$$s(\epsilon, \delta) \geq \frac{1}{2\epsilon^2} \log(2/\delta). \quad (1)$$

This bound implies that for any fixed δ , the number of samples required to achieve ϵ error scales polynomially in $1/\epsilon$. We call estimators satisfying this property *poly-precision* estimators. (We distinguish these from *exponential-precision* estimators, defined as estimators for which $s(\epsilon, \delta)$ scales logarithmically in $1/\epsilon$.)

Our central results are a classical algorithm that produces a poly precision estimate of a quantum circuit in

the above class, and a bound on the efficiency of this algorithm based on a measure of the circuit's negativity in a quasiprobability representation.

Quasiprobability representations.—A quasiprobability representation of a qudit over Λ is defined [11, 12] by a frame $\{F(\lambda) : \lambda \in \Lambda\}$ and a dual frame $\{G(\lambda) : \lambda \in \Lambda\}$, which are (generally over-complete) bases for the space of Hermitian operators acting on \mathbb{C}^d satisfying $A = \sum_{\lambda \in \Lambda} G(\lambda) \text{Tr}[AF(\lambda)]$ for all A . The space Λ can be continuous or discrete, and although many quasiprobability representations assume a phase space (symplectic) structure on Λ , this is not necessary. We can define quasiprobability distributions on Λ associated with a quantum state ρ , a unitary operator U and a measurement effect E to be

$$\begin{aligned} W_\rho(\lambda) &= \text{Tr}[F(\lambda)\rho], \\ W_U(\lambda'|\lambda) &= \text{Tr}(F(\lambda')UG(\lambda)U^\dagger), \\ W(E|\lambda) &= \text{Tr}[EG(\lambda)]. \end{aligned} \quad (2)$$

Tensor products of these dual frames gives a dual frame for the product space, and so these definitions extend in the obvious way from a tensor product $(\mathbb{C}^d)^{\otimes N}$ of N qudits to distributions on a phase space Λ^N .

The distribution $W_\rho(\lambda)$ is real-valued and satisfies $\sum_{\lambda \in \Lambda} W_\rho(\lambda) = 1$, much like a probability distribution, if the frame is normalized using $\sum_{\lambda \in \Lambda} F(\lambda) = I$. Similarly, the distributions $W_U(\lambda'|\lambda)$ and $W(E|\lambda)$ are normalized like corresponding conditional probabilities. The Born rule $\text{Pr}(E|\rho, U) = \text{Tr}(EU\rho U^\dagger)$, which gives the quantum probability for a measurement outcome given the state and process, is reproduced in the quasiprobability representation as would be expected in a probabilistic theory, by

$$\text{Pr}(E|\rho, U) = \sum_{\lambda, \lambda' \in \Lambda} W(E|\lambda') W_U(\lambda'|\lambda) W_\rho(\lambda). \quad (3)$$

This equation follows from the Born rule using the definition of the dual frames.

Importantly, the distributions of a quasiprobability representation will generally take on negative values, and so cannot be directly interpreted as probability distributions. The 1-norm of a quasiprobability distribution provides a natural measure of the amount of negativity, i.e., how much it deviates from a true probability distribution. We define the *negativity* \mathcal{M}_ρ of a state ρ as the 1-norm of its quasiprobability representation,

$$\mathcal{M}_\rho = \|W_\rho\|_1 = \sum_{\lambda \in \Lambda} |W_\rho(\lambda)|. \quad (4)$$

(The *mana* of a state using the discrete Wigner representation was introduced in Ref. [13] as a measure to bound the resources required for magic state distillation, and defined as the logarithm of the negativity used here.)

Analogously, we define the *negativity* \mathcal{M}_E of a measurement effect E to be

$$\mathcal{M}_E = \sum_{\lambda \in \Lambda} |W(E|\lambda)|, \quad (5)$$

and the *point-negativity* $\mathcal{M}_U(\lambda)$ and *negativity* \mathcal{M}_U of a unitary U to be

$$\mathcal{M}_U(\lambda) = \sum_{\lambda' \in \Lambda} |W_U(\lambda'|\lambda)|, \quad \mathcal{M}_U = \max_{\lambda \in \Lambda} \mathcal{M}_U(\lambda), \quad (6)$$

respectively. The negativities for states, unitaries and effects are lower-bounded by 1, 1 and $\text{Tr}(E)$ respectively, with equality if and only if the quasiprobability representation is nonnegative. These negativities will serve as a measure of the cost of each circuit element in our estimator.

Estimation procedure.—A quasiprobability representation provides an interpretation of the Born rule as the expectation value of a stochastic process. Specifically, if the quasiprobability representation of all elements in the circuit are nonnegative, then one may interpret the Born rule of Eq. (3) as the expected probability of the measurement outcome averaged over a set of trajectories through phase space [8, 9]. This stochastic interpretation no longer holds if the quasi-probability representation for any of the input states, gates, or measurements is negative. A standard perspective is that, in a quantum description, different trajectories in phase space can be assigned negative weights and can interfere with each other [10]. Monte Carlo sampling techniques may still be used, but the key problem is to identify an appropriate distribution to sample trajectories $\boldsymbol{\lambda} = (\lambda_0, \dots, \lambda_L)$ through phase space, where λ_0 is associated with the preparation and λ_l and λ_{l-1} are associated with the l th unitary. Eq. (3) becomes $\text{Pr}(E|\rho, U) = \sum_{\boldsymbol{\lambda}} W(\boldsymbol{\lambda})$ with

$$W(\boldsymbol{\lambda}) = W(E|\lambda_L) \left[\prod_{l=1}^L W_U(\lambda_l|\lambda_{l-1}) \right] W_\rho(\lambda_0). \quad (7)$$

Using an approach reminiscent of quantum Monte Carlo methods for fermion systems, we could sample trajectories from a true (nonnegative) probability distribution obtained from the absolute value of the quasiprobability, keeping track of the sign of the sampled trajectory. Consider the distribution of trajectories given by

$$\text{Pr}(A = \boldsymbol{\lambda}) = \frac{|W(\boldsymbol{\lambda})|}{\mathcal{M}_c} \quad (8)$$

where $\mathcal{M}_c = \sum_{\boldsymbol{\lambda}} |W(\boldsymbol{\lambda})|$ measures the negativity of the entire circuit, and we regard $\boldsymbol{\lambda}$ as a realization of a random variable A . An estimate based on a single realization $\boldsymbol{\lambda}$ is given by $\hat{q}_1(\boldsymbol{\lambda}) = \mathcal{M}_c \text{Sign}[W(\boldsymbol{\lambda})]$, where $\text{Sign}[\cdot] = \pm 1$ depending on the sign of the input. The expected value of this estimate gives the desired Born

rule probability

$$\begin{aligned} \langle \hat{q}_1(A) \rangle &= \sum_{\lambda} \hat{q}_1(\lambda) \Pr(A = \lambda) = \sum_{\lambda} \text{Sign}[W(\lambda)] |W(\lambda)| \\ &= \sum_{\lambda} W(\lambda) = \Pr(E|\rho, U). \end{aligned} \quad (9)$$

Note that this estimator minimizes the range of A [10], and so provides the best bound on the number of samples required to obtain a fixed precision using the Hoeffding inequality. Sampling from the distribution (8) is also the optimal estimator over the space of trajectories in that it has the smallest variance (see Supplementary Information). Unfortunately, without any additional structure, this estimator will in general be impractical for two reasons: there is no known efficient method to compute \mathcal{M}_c , and sampling from the distribution (8) will in general be inefficient in N .

To develop an efficient procedure, we sample trajectories λ following a Markov chain, using (true) probabilities and conditional probabilities at each timestep. Consider an input product state $\rho = \otimes_{n=1}^N \rho_n$, which has an efficient description $W_\rho(\lambda)$ in the quasi-probability representation which may be negative. We sample the initial point of the trajectory λ_0 from the modified distribution $\Pr(\lambda_0) = |W_\rho(\lambda_0)| / \mathcal{M}_\rho$. We construct a full trajectory λ by sampling λ_l at each timestep $l = 1, \dots, L$ in the circuit from the conditional distribution $\Pr(\lambda_l | \lambda_{l-1}) = |W_{U_l}(\lambda_l | \lambda_{l-1})| / \mathcal{M}_{U_l}(\lambda_{l-1})$ given by the unitary gate U_l . If the unitary U_l at each timestep l of the circuit has an efficient description in the quasi-probability representation, for example it consists of quantum gates acting on a fixed number of systems, then this distribution can be sampled efficiently. We note that this efficiency comes at a cost, as trajectories are no longer sampled from the optimal distribution (8).

An estimate \hat{p}_1 based on a single trajectory λ of our Markov chain protocol is given by

$$\begin{aligned} \hat{p}_1(\lambda) &= \mathcal{M}_\rho \text{Sign}[W_\rho(\lambda_0)] \\ &\times \prod_{l=1}^L \left[\mathcal{M}_{U_l}(\lambda_{l-1}) \text{Sign}[W_{U_l}(\lambda_l | \lambda_{l-1})] \right] W(E|\lambda_L). \end{aligned} \quad (10)$$

Unlike \hat{q} , this estimate is guaranteed to be an efficiently computable function of the sampled path λ . We note that \hat{p}_1 can lie outside the unit interval, but nonetheless gives an unbiased estimate of the Born rule probability, $\langle \hat{p}_1(A) \rangle = \Pr(E|\rho, U)$, precisely as in Eq. (9). Further, we note that \hat{p}_1 lies in the interval $[-\mathcal{M}_\rightarrow, +\mathcal{M}_\rightarrow]$, where we have defined \mathcal{M}_\rightarrow to be the *total forward negativity bound* of the circuit:

$$\mathcal{M}_\rightarrow = \mathcal{M}_\rho \prod_{l=1}^L \mathcal{M}_{U_l} \max_{\lambda_L} |W(E|\lambda_L)|. \quad (11)$$

Let \hat{p}_s be the average of \hat{p}_1 over s independent samples of λ . Using the boundedness and unbiasedness properties

of \hat{p}_1 , the Hoeffding inequality yields an upper bound on the rate of convergence of the average \hat{p}_s . Specifically, \hat{p}_s will be within ϵ of the quantum probability $\Pr(E|\rho, U)$ with probability $1 - \delta$ if a total of

$$s(\epsilon, \delta) = \frac{2}{\epsilon^2} \mathcal{M}_\rightarrow^2 \ln(2/\delta) \quad (12)$$

samples are taken. Consequently, if the total forward negativity bound \mathcal{M}_\rightarrow grows at most polynomially with N , then our protocol gives an efficient estimate \hat{p}_s of the quantum probability $\Pr(E|\rho, U)$ to within $\epsilon = 1/\text{poly}(N)$, with an exponentially small failure probability. That is, for circuits with a polynomially-bounded total forward negativity bound, \hat{p}_s is a poly-precision estimator of the Born rule probability and we can sample \hat{p}_s efficiently in N . We note that the total forward negativity bound \mathcal{M}_\rightarrow of (11) is insensitive to the measurement negativity \mathcal{M}_E , instead depending only on $\max_{\lambda_L} |W(E|\lambda_L)|$.

Any efficiently computable symmetry of the Born rule can be used to give a variant on the procedure defined above. The rate of convergence of the estimator need not be symmetric under these Born rule symmetries, and so such a variant may provide an advantage. Two examples of such symmetries – the time reversal symmetry that exchanges states and measurement effects in a unitary circuit, and the regrouping of unitaries into different elementary gates – are explored in the Supplementary Information. In particular, a variant procedure is presented for which the total negativity bound is insensitive to the negativity of the initial state \mathcal{M}_ρ .

Example: Estimation with the discrete Wigner function.—The odd- d qudit stabilizer subtheory and the associated discrete Wigner function provide a canonical example for demonstrating the use of our algorithm; see also Ref. [10]. Using this discrete Wigner representation for our estimation algorithm, the nonnegativity of the stabilizer subtheory [14, 15] ensures that stabilizer states, gates, and rank-1 measurements have negativity $\mathcal{M}_{\rho_n/U_l/E_n} = 1$ and so are “free” resources. Moreover, due to the existence of nonnegatively represented operations that are not in the stabilizer polytope [8], our approach is efficient on a strictly larger set of circuits than those of Ref. [1]. Circuits with operations possessing negativity strictly greater than 1, such as magic states and non-Clifford gates, can still be estimated but now come at a cost. Provided the total negativity bound grows at most polynomially in N , our protocol provides an efficient estimator.

As an example, consider a circuit with an input state given by a product state of k qutrit magic states $\frac{1}{\sqrt{3}}(|0\rangle + \xi|1\rangle + \xi^8|2\rangle)$, with $\xi = \exp(2\pi i/9)$, together with stabilizer $|0\rangle$ states in a 100-qutrit random Clifford circuit, and estimate the probability of measuring $|0\rangle$ on the first qutrit of the output. The total forward negativity bound of this circuit scales exponentially in k and consequently




FIG. 1: (Color online) Plot of the difference $\hat{p}_s - \langle \hat{p} \rangle$ between the estimated probability and the true probability of the outcome $|0\rangle\langle 0|$ for the first qutrit as a function of the number of magic states k . Each data point represents a random 100-qutrit Clifford circuit with the non-magic states initialized to the $|0\rangle$ state. The number of samples $s(k)$ was chosen using Eq. (12) with target precision $\epsilon = 0.01$ (indicated by the solid line) with 95% confidence ($\delta = 0.05$), so the number of samples increases exponentially with k (color scale).

the number of samples required to guarantee a fixed precision scales exponentially in k by Eq. (12). The results of our numerical simulations, shown in Fig. 1, indicate that our estimator does indeed converge with an appropriately chosen number of samples. Moreover, while the true precision of \hat{p}_s in our simulations is often orders of magnitude better than the target precision, there are circuits that come close to saturating the target precision, suggesting that our bound cannot be substantially improved without further detailed knowledge of the circuit.

Discussion.—Our results highlight the role of the total negativity of a circuit as a resource required for a quantum computer to outperform any classical computer. In particular, any circuit element that is represented non-negatively does not contribute to the total negativity bound and can be viewed as a “free” resource within the algorithm. Other circuit elements have an associated cost quantified by their negativity, *unless* they appear at the final timestep of the algorithm. This latter observation motivates us to exploit the time-reversal and other symmetries of the Born rule, seeking to minimize the total forward (or reverse) negativity bound. In particular, one could seek equivalent circuits wherein negative operations can be replaced with nonnegative ones by using negative initial states or measurements via gate teleportation. By choosing the forward or reverse procedure as appropriate, the efficiency can be made insensitive to the negativity of these initial states or measurements.

It also motivates us to identify quasiprobability representations in which many of the circuit elements of interest are represented nonnegatively. Interesting and

relevant examples abound, beyond the well-studied qudit discrete Wigner function. Discrete Wigner functions for qubits ($d = 2$) can be defined for which all stabilizer states with real coefficients (rebits) and all CSS-preserving unitaries are nonnegatively represented [16]. The range of quasiprobability representations introduced in Ref. [17] represent discrete subgroups of $U(2)$ on a single qubit nonnegatively, but have no nonnegative entangling gates; as such representations can represent certain non-stabilizer single-qubit states nonnegatively, they may be useful for estimation in circuits for gate synthesis. There is also flexibility in how a quasiprobability representation is defined. For example, a given quasiprobability representation can be modified to describe more states nonnegatively at the expense of a decreasing set of nonnegative measurements, and vice versa, by exploiting the structure of the dual frames. Overcomplete frames provides freedom in the choice of dual frame, and the negativity of unitaries and measurements will depend on this choice. As the dual frames formalism itself captures the relationship between quantum states and measurements, there is also freedom in the definition of the quasiprobability representation of unitaries beyond that given by Eq. (2). Finally, again using the freedom in the choice of dual for overcomplete frames, it is possible to switch between frames throughout a single circuit. These freedoms can be used to minimize the total negativity bound of the circuit, allowing more efficient estimators.

Our procedure can be applied to infinite-dimensional Hilbert spaces using any of the range of quasiprobability representations with continuous phase spaces developed in the study of quantum optics, by performing an appropriate discretization as in Ref. [5]. In this case, the negativity of distributions is quantified by integrating the absolute value of the distributions over the phase space, and is directly related to the *volume* of negativity [18]. We note that the resulting estimator can be applied to quantum optics experiments including states and measurements with negative Wigner function, such as photon number Fock states, and so may provide additional insight into the classical simulation cost of boson sampling [19]. While there exist means to efficiently estimate the outcome probability of a specific linear optics circuit with Fock state input and measurement [6], our estimation procedure extends these results by providing a general method for estimating outcome probabilities of such linear optical circuits for any input and output together with a bound on the efficiency of this estimation based on the volume of negativity of these states. In addition, our estimation can easily incorporate squeezing, as well as the loss and noise mechanisms common to linear optics experiments.

There are two natural ways to extend our results to circuits that include intermediate measurements and conditional operations based on them. First, one could replace the measurement and conditional operation with a

coherently-controlled operation, and delay the measurement to the end. We note that such controlled operations can be negative, even if the measurement and classically-controlled operation are both nonnegative. Second, our algorithm can be used to directly estimate the probabilities of the intermediate measurements and to sample from them. In this case, the required precision is exponential in the number of intermediate measurements in order to calculate conditional probabilities for subsequent use in the algorithm. Thus, in general, both approaches require resources that are exponential in the number of intermediate measurements.

Finally, our estimation procedure provides insight into the study of operationally meaningful measures of non-classical resources in quantum computation. Negativity in a quasiprobability representation has long been used as an indicator of quantum behaviour, but only recently has it been quantified as a resource for quantum computation [13]. Our results provide a related operational meaning of this resource: as a measure that bounds the efficiency of a classical estimation of probabilities.

The authors are grateful to J. Emerson, C. Ferrie, S. Flammia, R. Jozsa and A. Krishna for helpful discussions. This work is supported by the ARC via the Centre of Excellence in Engineered Quantum Systems (EQuS) project number CE110001013.

[1] S. Aaronson and D. Gottesman, Phys. Rev. A **70**, 052328 (2004).
 [2] L. G. Valiant, SIAM J. Comput. **31**, 1229 (2002).

[3] B. M. Terhal and D. P. DiVincenzo, Phys. Rev. A **65**, 032325 (2002).
 [4] S. D. Bartlett, B. C. Sanders, S. L. Braunstein, and K. Nemoto, Phys. Rev. Lett. **88**, 097904 (2002).
 [5] V. Veitch, N. Wiebe, C. Ferrie, and J. Emerson, New J. Phys. **15**, 013037 (2013).
 [6] L. Gurvits, in *Mathematical Foundations of Computer Science*, Lecture Notes in Computer Science **3618**, 447-458 (2005).
 [7] C. Gardiner and P. Zoller, *Quantum Noise: A Handbook of Markovian and Non-Markovian Quantum Stochastic Methods with Applications to Quantum Optics* (Springer-Verlag, Berlin, 3rd ed., 2004).
 [8] V. Veitch, C. Ferrie, D. Gross, and J. Emerson, New J. Phys. **14**, 113011 (2012).
 [9] A. Mari and J. Eisert, Phys. Rev. Lett. **109**, 230503 (2012).
 [10] D. Stahlke, Phys. Rev. A **90**, 022302 (2014).
 [11] C. Ferrie and J. Emerson, New J. Phys. **11**, 063040 (2009).
 [12] C. Ferrie, Rep. Prog. Phys. **74**, 116001 (2011).
 [13] V. Veitch, S. A. H. Mousavian, D. Gottesman and J. Emerson, New J. Phys. **16**, 013009 (2014).
 [14] C. Cormick, E. F. Galvao, D. Gottesman, J.-P. Paz, and A. O. Pittenger, Phys. Rev. A **73**, 012301 (2006).
 [15] D. Gross, J. Math. Phys. **47**, 122107 (2006).
 [16] N. Delfosse, P. A. Guerin, J. Bian, and R. Raussendorf, Phys. Rev. X **5**, 021003 (2015).
 [17] J. J. Wallman and S. D. Bartlett, Phys. Rev. A **85**, 062121 (2012).
 [18] A. Kenfack and K. Życzkowski, J. Opt. B: Quantum Semiclass. **6**, 396 (2004).
 [19] S. Aaronson and A. Arkhipov, Proc. ACM Symposium on Theory of Computing, San Jose, CA pp. 333-342 (2011).