Quantifying quantum speedups: improved classical simulation from tighter magic monotones

James R. Seddon, Bartosz Regula, Hakop Pashayan, Yingkai Ouyang, and Earl T. Campbell

Submitted to arXiv 14/02/2020, arXiv ID pending.

1 Overview

In the stabilizer circuit model of quantum computation, universality requires a resource known as magic. In standard architectures for fault-tolerant computation, pure magic states must be produced through the costly process of magic state distillation[1;2]. The Gottesman-Knill theorem shows that stabilizer circuits — comprised of Clifford gates, Pauli measurements and classical feed-forward — acting on stabilizer state inputs are efficiently classically simulable. A common strategy for simulation of non-Clifford circuits is to gadgetize them, so that they can be modelled as a sequence of stabilizer operations on a tensor product of magic states. However, these gadgetized circuits can only be classically simulated with a runtime overhead that is expected to be exponential in the number of magic states[3;4]. The resource theory of magic is a tool that can be used for quantifying the cost of classical simulation of general non-stabilizer circuits, as well as optimizing protocols for magic state distillation.

For odd dimension qudits, there exists an elegant and well-studied theory of stabilizer computation based on the discrete Wigner function[5–7], but many of the convenient properties of the formalism break down when applied to the important qubit case[8]. Alternative magic monotones for qubits include the stabilizer rank and extent[9–11], as well as the robustness of magic[12]. Techniques based on stabilizer rank and extent represent the state of the art in stabilizer-based classical simulation, but the associated monotones are defined only for pure states. Conversely, robustness of magic is defined for all mixed states on qubits, but the associated quasiprobability simulators are generally slower than stabilizer rank techniques. One of the key motivations for the current work is to develop a unified framework enabling us to leverage the power of stabilizer rank techniques, while remaining applicable to general mixed states on qubits.

In this paper, we study three new magic monotones for qubit mixed states, respectively called the mixed-state extent, dyadic negativity, and generalized robustness, all of which can be viewed as generalizations of the pure-state extent. In the first part of the paper, we give a series of results which apply in specific cases of interest. One key result is that the monotones are multiplicative for product states where each tensor factor is a single-qubit state, establishing our three quantifiers as the first qubit magic monotones which exhibit multiplicativity for general mixed states. We then use the monotones to benchmark the efficiency of state conversion under stabilizer protocols, establishing bounds applicable both to probabilistic one-shot protocols as well as to the asymptotic rates of transformation. In the second part of the paper, we show how each of our three monotones has a direct operational interpretation with respect to a classical simulation task. In particular, we show that the dyadic negativity corresponds to the runtime for a quasiprobability simulator[6], and that we can achieve a speedup exponential in the number of consumed magic states when comparing with previously known simulators of this type[12;13]. We also give an algorithm associated with the mixed-state extent, generalizing simulation techniques based on stabilizer rank[10;11] that were previously only defined for pure states. In addition, we prove that the runtime needed to achieve a fixed precision can be orders of magnitude smaller than previously thought. We furthermore show how our results can be extended to the general framework of quantum resource theories[14] leading to simulation algorithms beyond the theory of magic states, and present an explicit application in the resource theory of quantum coherence[15].

2 Preliminaries

Previous magic monotones.— We first review the previously known magic monotones: stabilizer rank, extent and robustness of magic. The exact **stabilizer rank**[9;10] χ is the smallest number of terms needed to write a given pure magic state $|\psi\rangle$ as a superposition of pure stabilizer states, that is: $\chi(\psi) = \min\{k : |\psi\rangle = \sum_{j=1}^{k} c_j |\phi_j\rangle\, , |\phi_j\rangle \in \mathcal{V}_n\}$ with $\mathcal{V}_n$ denoting the set of n-qubit pure stabilizer states. Roughly speaking, stabilizer rank simulators work by efficiently simulating each stabilizer branch of the superposition in turn, so that the runtime scales linearly with the number of terms k in the decomposition. The discontinuous character of stabilizer rank can make it difficult to work with analytically; it is not multiplicative and the best lower bounds for many-qubit states are very loose. In contrast, the pure-state **extent** ξ is a continuous monotone which is based on minimizing the sum of the absolute values of the coefficients in a stabilizer decomposition: $\xi(\psi) = \min\{(\sum_j |c_j|)^2 : |\psi\rangle = \sum_j c_j|\phi_j\rangle, |\phi_j\rangle \in \mathcal{V}_n\}$. A previously known sparsification lemma[11] entails that we can approximate ψ with a proxy state that has $k = O(\xi(\psi)/\delta^2)$ known stabilizer terms, and is δ-close to $|\psi\rangle$ in the trace norm. The **robustness of magic**[12] $\mathcal{R}$ minimizes the ℓ_1-norm over all quasiprobability distributions of a given mixed state ρ, formally $\mathcal{R}(\rho) := \min\{||q||_1 : \rho = \sum_j q_j|\phi_j\rangle\langle\phi_j|\}$, where ϕ_j are stabilizer states. Using a Monte Carlo algorithm[6;12], one can classically estimate Pauli expectation values and Born rule probabilities with δ additive error and a runtime $O(||q||_1^2/\delta^2)$, which for an optimal decomposition is $O(\mathcal{R}(\rho)^2/\delta^2)$.

New magic monotones.— We introduce three magic monotones, all of which can be regarded as mixed-state extensions of the stabilizer extent ξ. Firstly, using the usual convex roof extension[16] we define the **mixed-state extent** Ξ to be the quantity

$$\Xi(\rho) := \min\Big\{\sum_j p_j \xi(\Psi_j) \;:\; \rho = \sum_j p_j |\Psi_j\rangle\langle\Psi_j|;\; p_j \in \mathbb{R}^+;\; \sum_j p_j = 1\Big\}. \tag{1}$$

We also generalize the notion of a quasiprobability distribution using the **dyadic negativity** monotone

$$\Lambda(\rho) := \min\Big\{\|\alpha\|_1 \;:\; \rho = \sum_j \alpha_j |L_j\rangle\langle R_j|;\; |L_j\rangle, |R_j\rangle \in \mathcal{V}_n;\; \alpha_j \in \mathbb{C}\Big\}. \tag{2}$$

The name reflects the fact that each $|L_j\rangle\langle R_j|$ comprises of a pair of vectors, and so is a dyad. This can be understood as a relaxation of the robustness of magic[12], such that $\Lambda(\rho) \leq \mathcal{R}(\rho)$ for all ρ. Viewing this quantity as the primal solution of a convex optimization problem, it is useful to state the equivalent dual formulation[17] in terms of witness operators. Defining $\mathcal{W}$ to be the set of Hermitian operators such that $|\langle L|W|R\rangle| \leq 1$ for all $|L\rangle, |R\rangle \in \mathcal{V}_n$, by strong duality we can write $\Lambda(\rho) = \max\{\mathrm{Tr}[W\rho] : W \in \mathcal{W}\}$. This bring us to our last monotone of interest: the **generalized robustness** Λ^+, defined as

$$\Lambda^+(\rho) = \max\big\{\mathrm{Tr}[W\rho] : W \in \mathcal{W}; W \geq 0\big\}. \tag{3}$$

It was already shown that for general resource theories, we have for pure states that $\Lambda^+(|\Psi\rangle\langle\Psi|) = \Lambda(|\Psi\rangle\langle\Psi|) = \Xi(|\Psi\rangle\langle\Psi|) = \xi(\Psi)$[17], so all of the monotones are indeed mixed-state generalisations of the extent. For general mixed states ρ, we find that $\Lambda^+(\rho) \leq \Lambda(\rho) \leq \Xi(\rho)$. It is important to note that both Λ and Λ^+ are computable as convex optimization problems and can be evaluated in practice. Although Ξ is notoriously difficult to evaluate, we will simplify its computation in some important cases of interest.

3 Resource-theoretic results

We provide a complete solution for our monotones for products of single qubit states as follows:

Theorem *11*. *Let σ_j be single-qubit states. Then*

$$\Lambda(\otimes_j \sigma_j) = \Xi(\otimes_j \sigma_j) = \Lambda^+(\otimes_j \sigma_j) = \prod_j \Lambda(\sigma_j) = \prod_j \Lambda^+(\sigma_j) = \prod_j \Xi(\sigma_j). \tag{4}$$

The proof uses the notion of strong duality and first finds the optimal form of single qubit magic witnesses to solve the single-qubit problem (see **Thm 4**, **Lem 6** and **Lem 7** of the manuscript for details). We then borrow results from the extent literature[11] to show that these single-qubit witnesses form valid multi-qubit witnesses under the tensor product to complete the proof.

One consequence of Thm. 11 is that the hard optimization problem of the mixed-state extent Ξ is easily solved for product states. Furthermore, prior to this work, there were no known strict multiplicativity results for resource monotones for mixed states in qubit magic theory. For instance, Howard and Campbell[12] found that the robustness of magic can be strictly sub-multiplicative, $\mathcal{R}(\rho \otimes \rho) < \mathcal{R}(\rho)^2$ for all non-stabilizer ρ

considered, while a qubit-based phase-space robustness considered by Raussendorf *et al.*[18] can behave super-multiplicatively. We compare our monotones to the robustness of magic $\mathscr{R}$, with **Thm 13** showing our three monotones are exponentially smaller in magnitude than $\mathscr{R}(\rho)$: we can find constants α, β with $\alpha > \beta$ such that $2^{\beta n} = \Lambda(\rho^{\otimes n}) = \Lambda^{+}(\rho^{\otimes n}) = \Xi(\rho^{\otimes n})$ but for the robustness we have $2^{\alpha n} \leq \mathscr{R}(\rho^{\otimes n})$. For example, for the Hadamard $|H\rangle$ state we find $\alpha = 0.271553$ and $\beta = 0.228443$, so the size of the gap is significant. This has important consequences for the runtime of various classical simulation algorithms, as we will consider shortly.

As an immediate application of our monotones, we use them to establish several bounds on state conversion and specifically the distillation of magic states, both in the practical one-shot regime as well as in the asymptotic framework of information theory. In particular, we use the generalized robustness Λ^{+} to establish fundamental bounds on the number of copies required to transform m copies of a state ρ to k copies of some target state ψ, explicitly taking into consideration the success probability p of the protocol as well the incurred error ε in the fidelity to the target state. Compared with another recent bound of this type[19], our bounds can perform significantly better in practical error regimes (see **Fig 4** in the paper for details), thus constituting the best known benchmark for the efficiency of magic state distillation. Furthermore, the bounds can be extended to the asymptotic regime, leading to a bound on the achievable rate R of the transformation $\rho^{\otimes m} \to \psi^{\otimes Rm}$ as $m \to \infty$. Since the monotone Λ^{+} can be evaluated as a semidefinite program, this leads to computable bounds on the efficiency of distillation as well as the asymptotic rates of transformations in magic theory.

4 Classical simulation algorithms

Dyadic negativity simulator.— The dyadic negativity minimizes the negativity with respect to stabilizer dyads and complex-valued quasiprobabilities. Although the basis of such dyads contains non-Hermitian elements, we show that decompositions of this kind can give rise to a classical simulation algorithm:

Theorem ***14 (informal version)*** *Let an n-qubit initial state with known decomposition into dyads* $\rho = \sum_j \alpha_j |L_j\rangle\langle R_j|$ *where* $||\alpha||_1 = \Lambda(\rho)$*. Let* $\mathscr{E}$ *be a sequence of* T *stabilizer-preserving operations, each acting on a few qubits. Then given a stabilizer projector* Π*, we can estimate the Born rule probability* $\mu = \mathrm{Tr}(\Pi \mathscr{E}[\rho])$ *with probability* $1 - p_{\mathrm{fail}}$ *and additive error* ϵ *within a runtime*

$$\frac{\Lambda(\rho)^2}{\epsilon^2} \log(p_{\mathrm{fail}}^{-1}) \mathrm{poly}(n, T). \tag{5}$$

In addition to employing dyadic decompositions, the algorithm underpinning this result has several other novel modifications to standard Monte Carlo techniques. For instance, instead of using the trace to evaluate the probabilities of different trajectories, we instead use the trace-norm. Since **Thm 13** tells us Λ can be exponentially smaller than $\mathscr{R}$ and other qubit known measures of negativity, this gives an exponential improvement in runtime. For n Hadamard magic states, the runtime of quasiprobability simulation improves from $O(2^{0.7372n})$[20] to $O(2^{0.4569n})$.

Stabilizer rank simulator.— Simulation methods based on the stabilizer rank can be used for both Born rule probability estimation and approximately sampling from the output distribution of a quantum circuit. However, existing stabilizer rank simulation algorithms only apply to pure states. We extend these techniques to mixed states, with a runtime connected to the mixed-state extent Ξ. By switching to approximation of density matrices instead of pure state vectors, we quadratically improve on previous sparsification lemmas. This is captured in the following theorem.

Theorem ***15 (informal version)****. We can approximate a given pure state* $|\psi\rangle\langle\psi|$ *with a density matrix* ρ_1 *that is* δ *close in trace-norm. Moreover,* ρ_1 *is an ensemble of pure states each with stabilizer rank no more than* $4\xi(\psi)/\delta$*. There are some mild caveats that* δ *is not too small.*

This represents a factor of $1/\delta$ improvement (see **Fig 5** for numerical comparisons) to the previous sparsification lemma[11] which led to a stabilizer rank of $4\xi(\psi)/\delta^2$ with similar caveats. Combining our sparsification lemma with optimised mixed-state extent decompositions, we get:

Theorem 18 (informal version) *Let ρ be a state with known mixed-state extent decomposition. Then there is a classical algorithm that approximately samples from the probability distribution associated with a sequence of Pauli measurements on ρ. Our samples come from a distribution that is δ-close in ℓ_1-norm to the actual distribution, and each sample has an expected runtime $\mathbb{E}(T) = O(\Xi(\rho)/\delta^3)$, again with mild caveats on δ not being too small. If ρ is a product of single qubit states, every term in the decomposition can have the same pure-state extent, leading to zero variance in the runtime.*

There are two notable technical advances here: a factor $1/\delta$ improvement in runtime inherited from Thm 15; and the rather surprising outcome that sampling can often be performed without variance in the runtime.

Runtime comparison.— Given the wide selection of classical simulation algorithms, a natural question is what methods have the best runtimes. Our resource-theoretic results provide a rigourous foundation for such a comparison. The stabilizer rank simulator runtime scales with $\Xi(\rho)$ and the dyadic negativity simulator runtime scales with $\Lambda(\rho)^2$, so unless $\Lambda(\rho) \ll \Xi(\rho)$ the advantage lies with stabilizer rank methods. In general $\Lambda(\rho) \le \Xi(\rho)$, but Thm 11 shows that $\Xi(\rho) = \Lambda(\rho)$ for products of single-qubit states, so in this setting stabilizer rank simulators are likely preferred. However, stabilizer rank simulator have a worse dependence on δ and some additional constant factors, so there may still be some problems and parameter regimes for which a quasiprobability approach is preferred. Our technical manuscript describes a third classical simulation method connected to the Λ^+ monotone, but we omit it here for brevity.

Generalization to other resources.— The framework of quantum resource theories [14] provides a general setting for the investigation of various properties of quantum systems, with magic being an instance of such a resource. Crucially, we show that several of our results are not limited to magic theory, but can be extended to any quantum resource satisfying suitable assumptions on the classical simulability of the given resource theory's free states and operations. As an immediate application, we use our dyadic simulator in **Thm 14** to simulate free operations in the resource theory of coherence [15], leading not only to the first such simulation algorithm for this resource, but also giving an explicit operational interpretation to the ℓ_1-norm of coherence — a fundamental and widely used quantifier of coherence, but so far not connected directly to any operational task — in characterizing the runtime of such simulation.

References

[1] S. Bravyi and A. Kitaev, Physical Review A **71**, 022316 (2005).
[2] E. T. Campbell, B. M. Terhal, and C. Vuillot, Nature **549**, 172 (2017).
[3] D. Gottesman, Physical Review A **57**, 127 (1998).
[4] S. Aaronson and D. Gottesman, Physical Review A **70**, 052328 (2004).
[5] V. Veitch, S. H. Mousavian, D. Gottesman, and J. Emerson, New Journal of Physics **16**, 013009 (2014).
[6] H. Pashayan, J. J. Wallman, and S. D. Bartlett, Physical Review Letters **115**, 070501 (2015).
[7] X. Wang, M. M. Wilde, and Y. Su, arXiv preprint arXiv:1812.10145 (2018).
[8] S. Mansfield and E. Kashefi, Physical review letters **121**, 230401 (2018).
[9] S. Bravyi, G. Smith, and J. A. Smolin, Physical Review X **6**, 021043 (2016).
[10] S. Bravyi and D. Gosset, Physical Review Letters **116**, 250501 (2016).
[11] S. Bravyi, D. Browne, P. Calpin, E. Campbell, D. Gosset, and M. Howard, Quantum **3**, 181 (2019).
[12] M. Howard and E. Campbell, Physical Review Letters **118**, 090501 (2017).
[13] R. S. Bennink, E. M. Ferragut, T. S. Humble, J. A. Laska, J. J. Nutaro, M. G. Pleszkoch, and R. C. Pooser, Physical Review A **95**, 062337 (2017).
[14] E. Chitambar and G. Gour, Rev. Mod. Phys. **91**, 025001 (2019).
[15] T. Baumgratz, M. Cramer, and M. B. Plenio, Phys. Rev. Lett. **113**, 140401 (2014).
[16] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters, Phys. Rev. A **54**, 3824 (1996).
[17] B. Regula, Journal of Physics A: Mathematical and Theoretical **51**, 045303 (2017).
[18] R. Raussendorf, J. Bermejo-Vega, E. Tyhurst, C. Okay, and M. Zurel, arXiv preprint arXiv:1905.05374 (2019).
[19] K. Fang and Z.-W. Liu, (2019), arXiv:1909.02540 .
[20] M. Heinrich and D. Gross, Quantum **3**, 132 (2019).