FAST ESTIMATION OF OUTCOME PROBABILITIES FOR QUANTUM CIRCUITS

Hakop Pashayan[1], Oliver Reardon-Smith[2], Kamil Korzekwa[2], and Stephen D. Bartlett[3]

[1]Institute for Quantum Computing and Department of Combinatorics and Optimization, University of Waterloo, ON, N2L 3G1 Canada
[1]Perimeter Institute for Theoretical Physics, Waterloo, ON, N2L 2Y5 Canada
[2]Faculty of Physics, Astronomy and Applied Computer Science, Jagiellonian University, 30-348 Kraków, Poland
[3]Centre for Engineered Quantum Systems, School of Physics, The University of Sydney, Sydney, NSW 2006, Australia

Overview

Given a quantum circuit of the form shown in figure 1 we provide state of the art classical algorithms for estimating and computing the Born-rule probability $p = P(x)$.

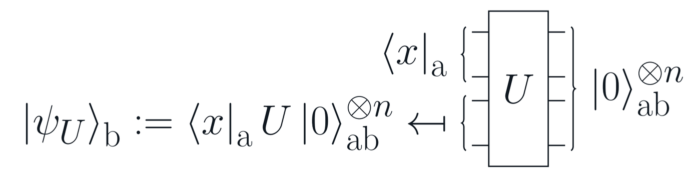

Fig. 1: (Read right to left) A unitary quantum circuit U acts on n qubits initially in the state $|0\rangle^{\otimes n}$, then $w \leq n$ qubits are measured in the computational basis and outcome x is observed with probability p.

Here we take "estimate" to mean that our algorithm takes as input a standard description of a quantum circuit with a measurement outcome with probability p, a failure probability δ and target precision ε and output a number $\hat{p}$ such that

$$P(|\hat{p} - p| > \varepsilon) < \delta. \tag{1}$$

Our agorithms are applicable where the unitary U consists of a (potentially large) number of basic Clifford gates

$$S \equiv \begin{bmatrix} 1 & 0 \\ 0 & i \end{bmatrix}, \quad H \equiv \frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \quad CX \equiv \begin{bmatrix} 1&0&0&0 \\ 0&1&0&0 \\ 0&0&0&1 \\ 0&0&1&0 \end{bmatrix}, \quad CZ \equiv \begin{bmatrix} 1&0&0&0 \\ 0&1&0&0 \\ 0&0&1&0 \\ 0&0&0&-1 \end{bmatrix}, \tag{2}$$

and a relatively small number of diagonal, single qubit non-Clifford gates

$$T_\phi \equiv \begin{bmatrix} 1 & 0 \\ 0 & e^{i\phi} \end{bmatrix}. \tag{3}$$

This gateset is universal for quantum computation even when we restrict to $\phi = \frac{\pi}{4}$ [3], however allowing general single qubit diagonal gates can dramatically improve performance. Circuits consisting of only Clifford gates applied to computational basis states are classically simulable in polynomial time [5], but are not even Turing complete [1].

Gadgetize and Compress

We employ a "reverse-gadgetization" of the T_ϕ gates, which may be summarised in the identity

$$\langle T_\phi^\dagger| \!\!-\!\!\oplus\!\!-\!\! |0\rangle \;=\; -\boxed{\tfrac{1}{\sqrt{2}}T_\phi}- , \tag{4}$$

where $|T_\phi^\dagger\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle + e^{-i\phi}|1\rangle\right)$. Gagetizing each T_ϕ gate gives

$$U|0\rangle_{\rm ab}^{\otimes n} = 2^{t/2} \langle T_{\boldsymbol\phi}^\dagger|_{\rm c} V |0\rangle_{\rm abc}^{\otimes n+t}. \tag{5}$$

The probability of observing outcome x is thus given by

$$p = 2^t \left\| \langle x|_{\rm a} \langle T_{\boldsymbol\phi}^\dagger|_{\rm c} V |0\rangle_{\rm abc}^{\otimes n+t} \right\|_2^2. \tag{6}$$

The Compress algorithm then operates on the stabilizer tableau of the gadgetized state $V|0\rangle_{\rm abc}^{\otimes n+t}$, its main purpose is to output the projector rank r the determinstic number v, a set of t-qubit stabilizer-generators $\{g_i \,|\, i = 1 \ldots t-r\}$ and a quantum circuit W such that

$$p = 2^{v-w} \langle T_{\boldsymbol\phi}^\dagger| \prod_{i=1}^{t-r} (I + g_i) |T_{\boldsymbol\phi}^\dagger\rangle, \tag{7}$$

$$= 2^{t-r+v-w} \left\| \langle 0|^{\otimes t-r} W |T_{\boldsymbol\phi}^\dagger\rangle \right\|_2^2, \tag{8}$$

note that where we previously had an $n+t$ qubit circuit V we now have a $t-r$ qubit circuit W. The number $r \leq \min\{t, n-w\}$ essentially counts the number of stabilizers that the Compress algorithm can delete from the stabilizer tableau. The gadegitzation and Compress algorithm take polynomial time.

Estimate

The Estimate algorithm repeatedly queries a subprocedure we call RawEstim. When provided with the output of the Compress algorithm and a pair of numbers s and L RawEstim produces an estimate $\hat{p}$ of the outcome probability p such that for all $\epsilon_{\rm tot} > 0$ and $\epsilon \in (0, \epsilon_{\rm tot})$:

$$\Pr(|\hat{p} - p| \geq \epsilon_{\rm tot}) \leq 2e^2 \exp\left(\frac{-s(\sqrt{p+\epsilon} - \sqrt{p})^2}{2(\sqrt{\xi^*} + \sqrt{p})^2}\right) + \exp\left(-\left(\frac{\epsilon_{\rm tot} - \epsilon}{p + \epsilon}\right)^2 L\right), \tag{9}$$

where ξ^* is the *stabilizer extent* of the state $|T_{\boldsymbol\phi}^\dagger\rangle$ which is multiplicative for single qubit states and hence for a fixed ϕ is exponential in the number of T_ϕ gates in the original circuit. RawEstim works by replacing the sum over 2^t stabilizer states appearing in equation (8) with a sample mean over s randomly sampled stabilizer states. We use a variant of Hoeffding's inequality to prove bounds on how far the sample mean can be from the true sum. We then use the fast-norm estimation algorithm of Ref. [4], which estimates the norm using the inner product with L randomly sampled equatorial states.
The run-time of the RawEstim algorithm scales as

$$\tau_{\text{RawEstim}} = O\left(st^2(t-r) + sLr^3\right). \tag{10}$$

From equation 9 we observe the s and L required to optain a particular ε,δ estimate depends on the unknown probability p we are estimating. Indeed one can show the s and L required are monotonically increasing in p. The Estimate algorithm therefore procedes by repeatedly upper bounding the target probability p by $p_k^* = \min\{1, \hat{p}_k + \varepsilon_k, p_{k-1}^*\}$, where ε_k is the error rate that attains the failure probability $\delta_k = \frac{6}{\pi^2 k^2}\delta_{\text{tot}}$ under the assumption that $p \leq p_{k-1}^*$. In order to attain increasingly tight upper bounds with only logarithmic overhead the Estimate algorithm chooses the paramaters s and L such that the run-time cost τ_{RawEstim} of calling RawEstim doubles each iteration. The choice of $\delta_k = \frac{6}{\pi^2 k^2}\delta_{\text{tot}}$ ensures that the total failure probability of the Estimate algorithm is less than $\sum_k \frac{6}{\pi^2 k^2}\delta_{\text{tot}} = \delta_{\text{tot}}$.

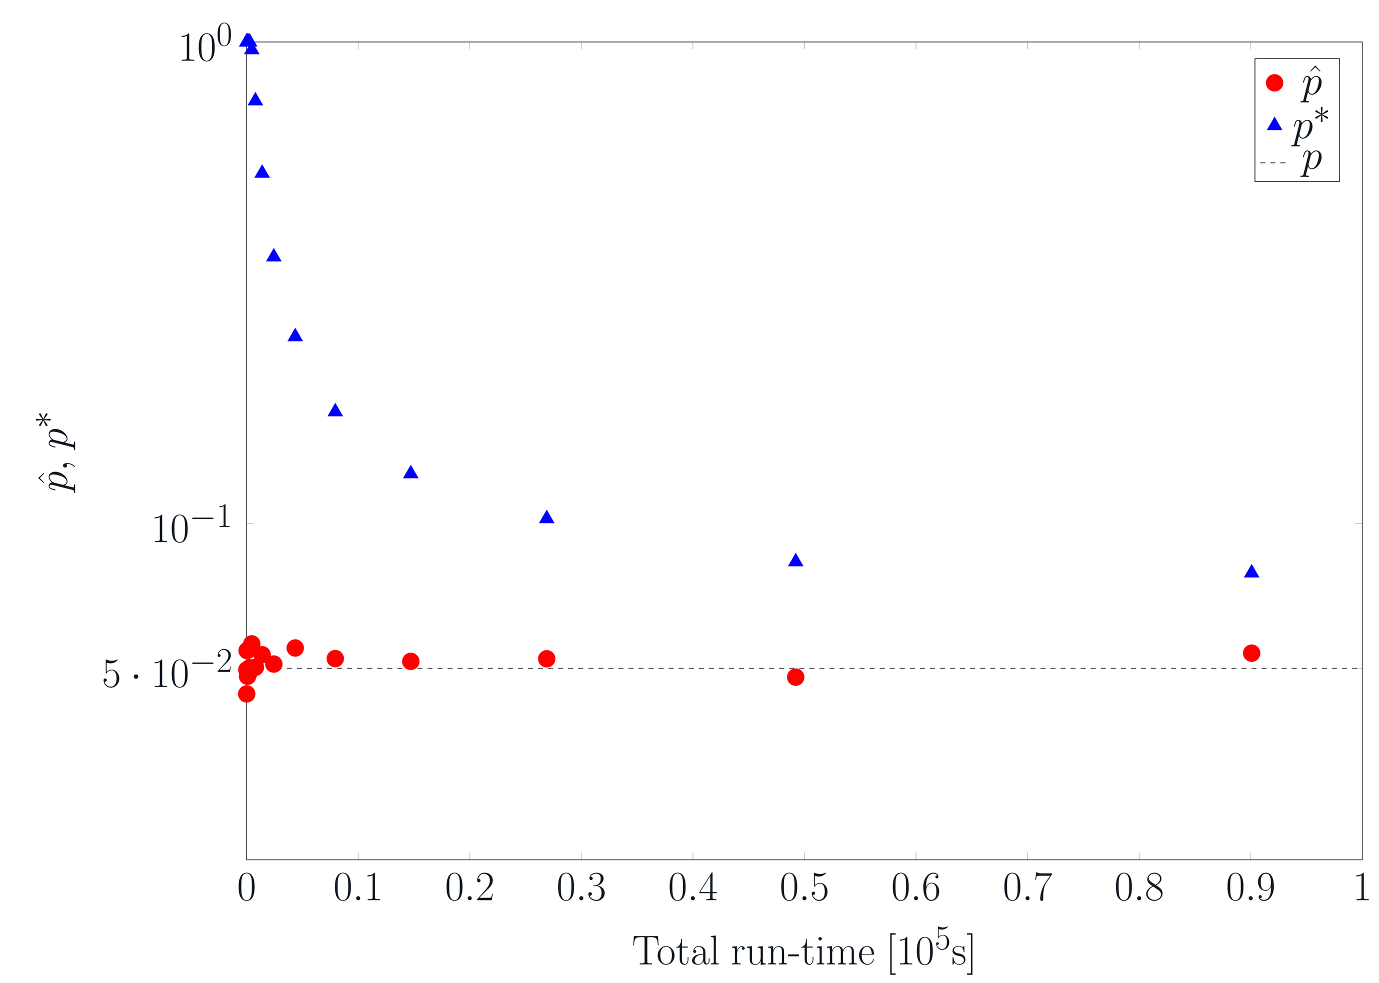

Fig. 2: Performance of the Estimate algorithm for an $n = 50$ qubit, $t = 60$ non-Clifford gate circuit constructed with a stabilizer extent $\xi^* \approx 3767$ equal to that of 52 $T_{\frac{\pi}{4}}$ gates. The Circuit consisted of 60 single qubit diagonal non-Clifford gates and approximately 2000 basic Clifford gates. The first $w = 8$ qubits were measured and the projector rank was $r = 10$. Each call to the RawEstim subprocedure provides an estimate $\hat{p}_k$ (red circles) and an upper bound $p_k^* = \hat{p}_k + \varepsilon_k$ (blue triangles) of the target Born rule probability p and incurs a time cost of $\tau_k = 2^k \tau_0$.

In order to improve the performance of our algorithm we extend the CH-form formalism of Ref. [4], providing an efficient algorithm to factorize CH-forms of tensor product states in and a "precomputation" method for computing the CH-form for a stabilizer state $UV|\psi\rangle$ directly from the CH-form of a state $U|\psi\rangle$ with runtime that is independent of the length of the circuit U. These improvements may be applicable more broadly.

Compute

The Compute algorithm directly computes the sum over 2^{t-r} terms appearing in equation (7), each term takes time $O(t(t-r))$ to calculate so the runtime of Compute is $O(2^{t-r}t(t-r))$. With high probability, r is very close to $\min\{t, n-w\}$ for random circuits with many Clifford gates, where w is the number of measured qubits. In figure 3 we compare the average runtime for the Compute algorithm to the average runtime for Qiskit's statevector simulator [2] for a number of randomly generated Clifford+T circuits. Our results show that despite its simplicity Compute can be a useful tool Compute can be effective in surprisingly challenging parameter regimes.

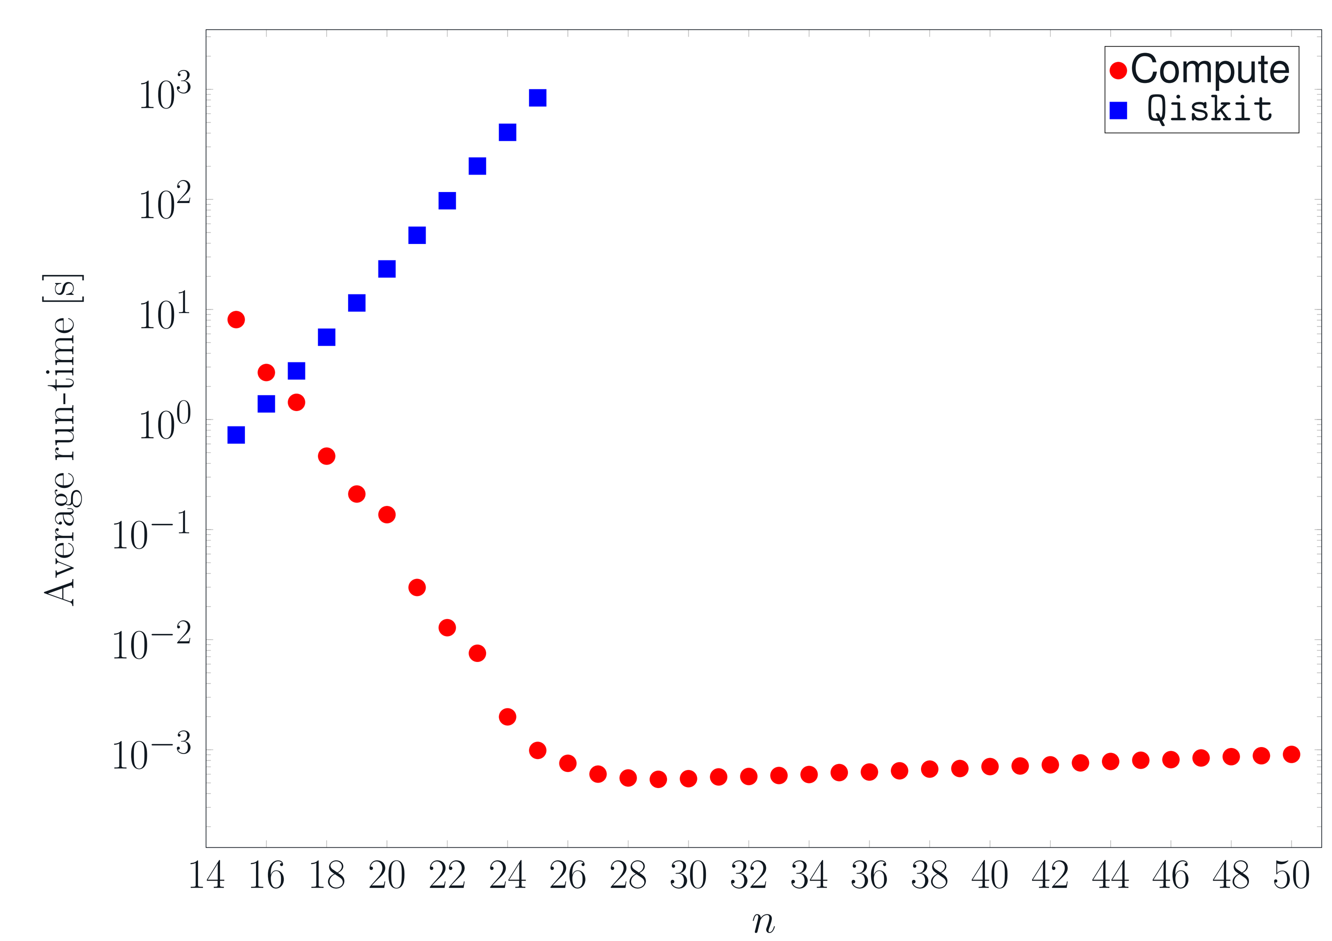

Fig. 3: Performance of the Compute algorithm for random circuits. These circuits were generated with 10^5 basic Clifford gates, 80 $T_{\frac{\pi}{4}}$ gates and 20 measured qubits. The averages are over 100 sampled circuit with the same circuits given to Qiskit's statevector simulator and the Compute algorithm. The runtime of Compute is exponential in $t-r$ where $r \approx n - w$ while the runtime of the statevector simulator is exponential in n.

Links

- Paper on arXiv - https://arxiv.org/abs/2101.12223
- Code on github - https://github.com/or1426/Clifford-T-estimator

References

[1] Scott Aaronson and Daniel Gottesman. "Improved simulation of stabilizer circuits". In: *Phys. Rev. A* 70 (5 Nov. 2004), p. 052328. DOI: 10.1103/PhysRevA.70.052328. URL: https://link.aps.org/doi/10.1103/PhysRevA.70.052328.

[2] Héctor Abraham et al. *Qiskit: An Open-source Framework for Quantum Computing.* 2019. DOI: 10.5281/zenodo.2562110.

[3] P. O. Boykin et al. "On universal and fault-tolerant quantum computing: a novel basis and a new constructive proof of universality for Shor's basis". In: *40th Annual Symposium on Foundations of Computer Science (Cat. No.99CB37039).* 1999, pp. 486–494. DOI: 10.1109/SFFCS.1999.814621.

[4] Sergey Bravyi et al. "Simulation of quantum circuits by low-rank stabilizer decompositions". In: *Quantum* 3 (2019), p. 181. DOI: 10.22331/q-2019-09-02-181.

[5] Daniel Gottesman. "The Heisenberg representation of quantum computers". In: *arXiv quant-ph/9807006* (1998). URL: https://arxiv.org/abs/quant-ph/9807006v1.