



HAL
open science

Classical simulation of quantum circuits with partial and graphical stabiliser decompositions

Aleks Kissinger, Renaud Vilmart, John van de Wetering

► **To cite this version:**

Aleks Kissinger, Renaud Vilmart, John van de Wetering. Classical simulation of quantum circuits with partial and graphical stabiliser decompositions. 17th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2022), Jul 2022, Urbana-Champaign, United States. pp.5:1 - 5:13, 10.4230/LIPIcs.TQC.2022.5 . hal-03606226

HAL Id: hal-03606226



<https://hal.science/hal-03606226v1>

Submitted on 11 Mar 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Classical simulation of quantum circuits with partial and graphical stabiliser decompositions

Aleks Kissinger   

University of Oxford, United Kingdom

John van de Wetering   

Radboud University Nijmegen, The Netherlands

University of Oxford, United Kingdom

Renaud Vilmart   

Université Paris-Saclay, CNRS, ENS Paris-Saclay, Inria, Laboratoire Méthodes Formelles, 91190, Gif-sur-Yvette, France

Abstract

Recent developments in classical simulation of quantum circuits make use of clever decompositions of chunks of magic states into sums of efficiently simulable stabiliser states. We show here how, by considering certain non-stabiliser entangled states which have more favourable decompositions, we can speed up these simulations. This is made possible by using the ZX-calculus, which allows us to easily find instances of these entangled states in the simplified diagram representing the quantum circuit to be simulated. We additionally find a new technique of *partial* stabiliser decompositions that allow us to trade magic states for stabiliser terms. With this technique we require only $2^{\alpha t}$ stabiliser terms, where $\alpha \approx 0.396$, to simulate a circuit with T-count t . This matches the α found by Qassim et al. [16], but whereas they only get this scaling in the asymptotic limit, ours applies for a circuit of any size. Our method builds upon a recently proposed scheme for simulation combining stabiliser decompositions and optimisation strategies implemented in the software QuiZX [15]. With our techniques we manage to reliably simulate 50-qubit 1400 T-count hidden shift circuits in a couple of minutes on a consumer laptop.

2012 ACM Subject Classification Theory of computation \rightarrow Quantum computation theory

Keywords and phrases ZX-calculus, Stabiliser Rank, Quantum Simulation

Funding *Aleks Kissinger*: Acknowledges support from AFOSR grant FA2386-18-1-4028.

John van de Wetering: NWO Rubicon Personal Grant

Acknowledgements JvdW and AK would like to thank James Hefford for useful discussions regarding decompositions of cat states in the ZX-calculus.

1 Introduction

A landmark result in the study of quantum circuit simulation is the Gottesman-Knill theorem [1], which states that a quantum circuit consisting of Clifford gates and stabiliser state inputs can be efficiently classically simulated. While Clifford gates are not universal for quantum computation, the Clifford+ T gate set, where we also allow the single qubit T gate, can approximate any qubit unitary to arbitrary precision [2, 17]. It is widely believed that the Clifford+ T gate set requires exponentially large classical resources to simulate in general. However, in practice certain methods exist for simulating Clifford+ T circuits significantly larger than one could ever hope to simulate directly, e.g. by state-vector calculation.

A particularly effective technique for circuits with relatively low numbers of non-Clifford gates is simulation by stabiliser decomposition. That is, general states (typically prepared via a Clifford+ T circuit) are first decomposed into weighted sums of stabiliser terms, each of which can be simulated efficiently via Gottesman-Knill. The *stabiliser rank* of a state $|\psi\rangle$ is the number of terms in the smallest possible decomposition of $|\psi\rangle$ as a weighted sum of

stabiliser states. While computing the exact stabiliser rank of a given state is expected to be a hard problem, various heuristics exist for obtaining stabiliser decompositions which scale with the number of non-Clifford gates in a circuit. By using magic state injection, we can write a state $|\psi\rangle$ obtained from a Clifford+T circuit as a Clifford circuit that takes as input a T magic state $|T\rangle := \frac{|0\rangle + e^{i\frac{\pi}{4}}|1\rangle}{\sqrt{2}}$ for each T gate in the original circuit. One can then decompose these T magic states, and hence $|\psi\rangle$ into a sum of stabiliser states. Since each of the individual terms can be simulated efficiently, the simulation cost of the state scales with the number of terms in the decomposition.

Naïvely, one can decompose each of the T magic states individually as a linear combination of the stabiliser states $|0\rangle$ and $|1\rangle$, obtaining 2^t terms. However, by decomposing larger non-stabiliser states at once, one can obtain lower simulation costs [7, 16]. Until now, such techniques have focused on decomposing many identical copies of a fixed magic state. For example, much work has gone into finding efficient decompositions of $|T\rangle^{\otimes t}$ for different values of t , with the best known decomposition scaling asymptotically as $2^{\alpha t}$ with $\alpha \approx 0.396$ [16].

In previous work by some of the current authors [15], a new method for simulation with stabiliser decompositions was introduced, which works by representing the circuit to be simulated as a *ZX-diagram* [8, 9], and then interleaving the decompositions of [7] with the diagram simplification strategy of [14] to reduce the number of stabiliser terms required, sometimes by many orders of magnitude. The scheme goes as follows:

1. Translate the circuit, together with the desired input state and measurement effect, to a ZX-diagram.
2. Simplify the diagram as much as possible using the rules of the ZX-calculus.
3. Pick a set of non-Clifford generators and decompose them, obtaining a weighted sum of diagrams with fewer non-Clifford generators.
4. Apply the previous two steps recursively to each of the diagrams until no non-Clifford generators remain, in which case each diagram is simplified to a single complex number.
5. The sum of these numbers gives the overall amplitude.

This allowed for the simulation of significantly larger circuits than before. While the previous best simulated a 50-qubit circuit with 64 T gates, in [15] they simulated 50-qubit circuits with up to 1400 T gates.

In this article we build on the work of both [16] and [15], by significantly improving Step 3. Our main contribution is to show that it is useful to consider different states than just $|T\rangle^{\otimes t}$ for decomposition. This leads to two improvements. The first is that if certain entangled states are present in the ZX-diagram to be simulated, then we can use a decomposition that asymptotically requires significantly fewer terms, having exponential parameter $\alpha \approx 0.25$ instead of $\alpha \approx 0.396$. While these special states need not appear in a generic ZX-diagram obtained from a Clifford+T circuit, our empirical data suggests that they often will. It is crucial here that we represent the intermediate stages of the simulation as a ZX-diagram, as opposed to a circuit with magic states as input, as the graph structure makes it possible to find the right entangled states to be decomposed.


Our second improvement is that one of these special decompositions can be adapted to lead to a decomposition of $|T\rangle^{\otimes 5}$ that uses only 3 terms but leaves one $|T\rangle$ in the resulting reduced terms, which effectively gives us a ‘4-to-3’ decomposition meaning we require $2^{\alpha t}$ terms with $\alpha \approx 0.396$ in the worst case. This matches the asymptotic upper bound reported in Ref. [16], however unlike that paper, we obtain this bound at a fixed finite size rather than in the limit.

We implemented our new decompositions in `quZX`, the software used in [15], and assessed the performance of our method by benchmarking the same quantum circuits as [15]. We


deal with explicitly when possible. Any diagram with 0 inputs and 0 outputs represents a scalar, and when it can be efficiently computed, we will simply do this, instead of leaving it as a diagram.

An important subclass of ZX-diagrams are the *Clifford ZX-diagrams*, i.e. those whose spiders only have phases equal to integer multiples of $\pi/2$. A Clifford ZX-diagram with no input wires always represents a stabiliser state, and conversely any stabiliser state can always be written as a Clifford ZX-diagram [3].

A convenient feature of ZX-diagrams is that the linear map they represent only depends on the connectivity of the diagram and not the actual positions of spiders or the direction of wires between them. This means in particular that we can topologically deform a diagram however we wish without changing its interpretation. This allows us to treat a ZX-diagram as an undirected open (multi-)graph, whose vertices are Z-spiders, X-spiders and H-gates. As H-gates are binary, it can be convenient to internalise them on their respective wires, and consider a new wire type which we call an *H-edge*, represented by a dashed blue line:



In addition to topological deformation, there are graphical rewrite rules that also preserve the semantics of the diagram, and in fact, *complete* axiomatisations of the ZX-calculus exist (e.g. [12, 13, 19]), which entirely capture the semantical equivalence of diagrams. We will not give a full such axiomatisation here, but only those rules we will need in this paper:



which by duality remain true when Z- and X-spiders are interchanged (i.e. when the colours are swapped).

Using these rules, it is possible to turn any diagram in a form where i) all spiders are Z-spiders and ii) all edges are H-edges (but inputs and outputs remain plain wires). A diagram in such a form is called *graph-like* [11]. These are the diagrams we will work with in this paper.

While the rewrite rules are unidirectional, we can apply one-way rewrite strategies that try to reduce some metric of the diagram, such as the number of spiders in a diagram. In particular, in [11] a simplification strategy was introduced that can remove most Clifford spiders, i.e. those spiders whose phase is an integer multiple of $\pi/2$, from a ZX-diagram. It is this rewrite strategy (or more specifically, the improved version from [14]) that was used in [15] to help in simplifying diagrams for simulation with stabiliser decompositions. As the details are not relevant to us, we will only describe some of the properties that the resulting diagrams enjoy, as these will be important for our results. We will call diagrams simplified by the rewrite strategy of [14] *reduced*. Note that a spider is called *internal* if it is not directly connected to an input or an output of the diagram.

► **Lemma 1.** *In reduced diagrams, no two Clifford spiders are neighbours.*


In other words, patterns like won't appear in the diagram.

► **Lemma 2.** *In reduced diagrams, no internal Clifford spider can have arity ≤ 2 .*

In other words, patterns like and won't appear.

► **Lemma 3.** *In reduced diagrams, no internal spider phases are odd multiples of $\frac{\pi}{2}$.*

In other words, patterns like $\vdots \textcircled{(2k+1)\frac{\pi}{2}} \vdots$ won't appear.




3 Stabiliser Rank and stabiliser decompositions

Since stabiliser states/operators are efficiently classically simulable, it is worthwhile to try to find decompositions of arbitrary quantum states as linear combination of stabiliser states with a small number of terms.

More formally, let $|\psi\rangle$ be an arbitrary state. Then a *stabiliser decomposition* of $|\psi\rangle$ is a decomposition $|\psi\rangle = \sum_k^n \lambda_k |\psi_k\rangle$ where $|\psi_k\rangle$ are all stabiliser states and $\lambda_k \in \mathbb{C}$ are some scalars. The smallest n for which such a decomposition exists is called the stabiliser rank of $|\psi\rangle$, and is denoted by $\chi(|\psi\rangle)$. For simulation purposes, we are obviously not only interested in $\chi(|\psi\rangle)$ but also in an associated decomposition.


In the Clifford+T case, most of the results on stabiliser decompositions have focussed on tensor products of the magic state $|T\rangle := \frac{1}{\sqrt{2}}(|0\rangle + e^{i\frac{\pi}{4}}|1\rangle)$, which we can represent as a ZX-diagram (up to normalisation) by $\textcircled{\frac{\pi}{4}}$. Whenever we have a spider with phase an odd multiple of $\pi/4$ we can unfuse a magic state:



A single magic state can be decomposed as a sum of two stabiliser states:

$$\textcircled{\frac{\pi}{4}} = \frac{1}{\sqrt{2}} (\textcircled{\frac{\pi}{2}} + e^{i\frac{\pi}{4}} \textcircled{\pi})$$

A naive decomposition of $\textcircled{\frac{\pi}{4}}^{\otimes t}$ would then be to apply this decomposition for all magic states, which would result in a decomposition with 2^t terms. More sophisticated decompositions exist however, such as:



We can use this decomposition to reduce T magic states pairwise, bringing the total number of terms down to $2^{\frac{t}{2}}$. There is also a $|T\rangle^{\otimes 6}$ decomposition requiring 7 terms, described in [7] and used in [15]. More recently, an improved decomposition of $|T\rangle^{\otimes 6}$ was found that only requires 6 terms [16]. To compare different decompositions it will be helpful to consider how many stabiliser terms would be needed if we were to decompose every non-Clifford phase in the diagram using that decomposition. If a decomposition reduces the T-count by r using p terms, we would need $p^{t/r} = 2^{\alpha t}$ stabiliser terms. This $\alpha = \log_2(p)/r$ will be the metric we want to minimise. For instance, the decomposition of [7] gives $\alpha \approx 0.468$ ($r = 6, p = 7$).

In the following, we will consider some entangled non-stabiliser states of which we can find a better decomposition than would be suggested by the number of magic states needed to write down these states.

4 Results

We present two different improvements to the stabiliser decompositions used before: using decompositions of certain entangled states instead of product states, and using ‘partial’

stabiliser decompositions where the terms are not Clifford, but merely contain a lower amount of magic states.

4.1 Cat states in the ZX-calculus

The states we are interested in here were introduced in Ref. [16] as *cat states*:

$$|\text{cat}_n\rangle := \frac{1}{\sqrt{2}}(\mathbb{I}^{\otimes n} + Z^{\otimes n})|T\rangle^{\otimes n} = \frac{1}{\sqrt{2^{n+1}}} \left(\textcircled{\frac{\pi}{4}}^{\otimes n} + \textcircled{\frac{5\pi}{4}}^{\otimes n} \right)$$

In Ref. [16], these states are used because i) they have a relatively small stabiliser rank and ii) because if the stabiliser rank of $|\text{cat}_n\rangle$ is bounded by c , then that of $|T\rangle^{\otimes n}$ is bounded by $2c$. For instance, they find a decomposition of $|\text{cat}_6\rangle$ using 3 stabiliser states, which hence gives a decomposition of $|T\rangle^{\otimes 6}$ using 6 stabiliser states. This improves upon the 6-to-7 decomposition of Ref. [7], and improves $\alpha \approx 0.468$ to $\alpha \approx 0.431$.

In this paper we focus on the use of the cat-states directly for stabiliser decompositions, rather than first transforming them into decompositions of sets of T states. First, notice that they can be easily expressed as ZX-diagrams:

$$|\text{cat}_n\rangle = \frac{1}{\sqrt{2}} \begin{array}{c} \textcircled{\frac{\pi}{4}} \\ \textcircled{\frac{\pi}{4}} \\ \vdots \\ \textcircled{\frac{\pi}{4}} \end{array}$$

We can then translate to ZX the decomposition of $|\text{cat}_6\rangle$ found in Ref. [16]¹:

$$\begin{array}{c} \textcircled{\frac{\pi}{4}} \\ \textcircled{\frac{\pi}{4}} \\ \textcircled{\frac{\pi}{4}} \\ \textcircled{\frac{\pi}{4}} \\ \textcircled{\frac{\pi}{4}} \\ \textcircled{\frac{\pi}{4}} \end{array} = \frac{1}{2} \begin{array}{c} \textcircled{-\frac{\pi}{2}} \end{array} + \frac{ie^{i\pi/4}}{\sqrt{2}} \begin{array}{c} \textcircled{\phantom{\frac{\pi}{4}}} \\ \textcircled{\phantom{\frac{\pi}{4}}} \\ \textcircled{\phantom{\frac{\pi}{4}}} \\ \textcircled{\phantom{\frac{\pi}{4}}} \\ \textcircled{\phantom{\frac{\pi}{4}}} \\ \textcircled{\phantom{\frac{\pi}{4}}} \end{array} - \frac{e^{i\pi/4}}{\sqrt{2}} \begin{array}{c} \textcircled{\frac{\pi}{2}} \\ \textcircled{\frac{\pi}{2}} \\ \textcircled{\frac{\pi}{2}} \\ \textcircled{\frac{\pi}{2}} \\ \textcircled{\frac{\pi}{2}} \\ \textcircled{\frac{\pi}{2}} \end{array}$$

Hence, if our diagram contained enough subdiagrams that look like $|\text{cat}_6\rangle$ to allow us to decompose all non-Clifford spiders using this decomposition we would get $3^{\frac{t}{6}} = 2^{\alpha t}$ terms, where $\alpha \approx 0.264$. We can in fact do even better by considering the decomposition of $|\text{cat}_4\rangle$, which has a stabiliser rank of only 2:

$$\begin{array}{c} \textcircled{\frac{\pi}{4}} \\ \textcircled{\frac{\pi}{4}} \\ \textcircled{\frac{\pi}{4}} \\ \textcircled{\frac{\pi}{4}} \end{array} = \frac{e^{-i\pi/4}}{\sqrt{2}} \begin{array}{c} \textcircled{-\frac{\pi}{2}} \end{array} + i \begin{array}{c} \textcircled{\phantom{\frac{\pi}{4}}} \\ \textcircled{\phantom{\frac{\pi}{4}}} \\ \textcircled{\phantom{\frac{\pi}{4}}} \\ \textcircled{\phantom{\frac{\pi}{4}}} \end{array}$$


If we could use only this decomposition we would get $2^{0.25t}$ terms.

The existence of these beneficial decompositions suggests a new strategy: at every step of the decomposition of the complete diagram, we look at every internal spider in the graph and the shape of its neighbourhood and we pick a target for decomposition that requires the

¹ The description of the last diagram in the decomposition is not the one given in [16]. However, the two can be shown to be equivalent using the ZX-calculus.

fewest number of terms per magic state. After each decomposition, we use a ZX-diagram simplification strategy (like the one in [15]) and then we repeat the procedure.


We cannot hope to always have an occurrence of $|\text{cat}_4\rangle$ or $|\text{cat}_6\rangle$ at hand. However, notice that we can always infer a decomposition of $|\text{cat}_n\rangle$ from one of $|\text{cat}_{n+k}\rangle$ ($k > 0$), by $|\text{cat}_n\rangle = (\langle 0|^{\otimes k} \otimes \mathbb{I}^{\otimes n})|\text{cat}_{n+k}\rangle$. This observation gives us for instance a decomposition of $|\text{cat}_3\rangle$ and $|\text{cat}_5\rangle$:




These two decompositions give us respectively $\alpha = 1/3 \approx 0.333$ and $\alpha \approx 0.317$.

Note that by Lemma 2, there will never be any occurrence of $|\text{cat}_2\rangle$ or $|\text{cat}_1\rangle$ in our simplified diagrams so that we now have a strategy for decomposing any $|\text{cat}_n\rangle$ for $n \leq 6$.

These decompositions apply when the ‘‘central’’ Clifford spider has phase 0. By Lemma 3 any internal Clifford spider will have phase 0 or π . The π case is easily taken care of by pushing the phase through one of the neighbours as follows:



Note that in practice, in our reduced diagrams the only occurrences of cat-states will be as *phase gadgets* [14, 10], which relate to cat states in the following way:




Hence, an n -legged phase gadget corresponds to a $|\text{cat}_{n+1}\rangle$ state, so that it can be decomposed more efficiently.

In Ref. [16] they also describe ways to get efficient stabiliser decompositions of $|\text{cat}_n\rangle$ for bigger n . For instance, they find a decomposition of $|\text{cat}_{10}\rangle$ which has 9 terms, which gives $\alpha \approx 0.317$. For simplicity of implementation, we don’t consider these decompositions of bigger states here.

4.2 Partial magic state decompositions

If we cannot find cat states in our diagram (which means that it only has non-Clifford spiders or that every Clifford spider has arity larger than 6), then we could fall back to the magic state decompositions described in Ref. [16]. Namely, as is observed in Ref. [16], we can


construct decompositions of bigger cat-states by the following identity:



Keep in mind that $|\text{cat}_n\rangle = (\langle 0|^{\otimes k} \otimes \mathbb{I}^{\otimes n})|\text{cat}_{n+k}\rangle$ so that we also get decompositions for n that are not equal to $4k + 2$. Using this construction of $|\text{cat}_{4k+2}\rangle$, we get a decomposition of $|T\rangle^{\otimes 4k+2}$ that uses $2 \cdot 3^k$ terms. This decomposition requires $2^{\alpha(4k+2)}$ terms where

$$\alpha = \frac{1 + k \log_2(3)}{4k + 2} \xrightarrow{k \rightarrow \infty} \frac{\log_2(3)}{4} \approx 0.396.$$

But as we will see now, we can in fact reach this asymptotic α without needing to decompose all magic states at once. The idea here is to use the decomposition of $|\text{cat}_6\rangle$ to reduce the T-count of the diagram by 4, using only 3 terms. This can be done whenever we have a T-count ≥ 5 , in the following way:



This is a (non-stabiliser) decomposition of the 5-qubit magic state which leaves one T-spider in each term. So each term effectively loses 4 magic states, so that we can view this as a 4-to-3 strategy. We get for this operation $\alpha = \frac{\log_2(3)}{4} \approx 0.396$ so that we reach this number concretely that in Ref. [16] was only reached asymptotically.

Additional advantages of this decomposition over the strategy using the $|\text{cat}_{4k+2}\rangle$ decomposition, are that i) it is simpler to handle and implement (as a smaller part of the diagram is involved) and ii) it more easily allows for the ZX-diagram reduction strategy of [15] to be used in-between rounds of decompositions, so that there is more possibility for reduction of the number of terms.

We remark that this type of ‘partial decomposition’ can be used to decompose any $|T\rangle^{\otimes 4k+1}$ states using the $|\text{cat}_{4k+2}\rangle$ decomposition, but doing so always results in the same α , whenever $k \geq 1$, so that there is no benefit to doing so (at least from an asymptotic perspective).

4.3 The full decomposition strategy

The conclusion of the previous sections is that some subdiagrams are better to decompose into stabilisers than others. In particular, in our diagrams we should look, in order of importance, for:

1. a Clifford spider with 4 neighbours ($\alpha = 0.25$),
2. a Clifford spider with 6 neighbours ($\alpha \approx 0.264$),
3. a Clifford spider with 5 neighbours ($\alpha \approx 0.317$),
4. a Clifford spider with 3 neighbours ($\alpha = 1/3 \approx 0.333$),

5. any 5 T-spiders ($\alpha \approx 0.396$).

Once the best subdiagram to decompose has been found, we apply the associated decomposition, then simplify each resulting diagram using the optimisation strategy of [14] as described in [15], and repeat the procedure for each term until we are left with diagrams that have fewer than 5 T-spiders and no cat-states, in which case we simply fall back to the usual magic-state decompositions.

5 Benchmarking

To assess the effectiveness of our approach, we performed simulations of the two families of circuits that were benchmarked in [15], and compare the associated running times with theirs. For a better comparison, we turned off parallelisation in most experiments. All subsequent tests were performed on a consumer laptop, equipped with an Intel Core i5-10400H CPU, 2.60GHz. By using a depth-first decomposition approach the amount of memory needed for these benchmarks was insignificant. In these benchmarks we refer to the approach of [15] as *BSS*, which stands for Bravyi, Smith, Smolin, as it uses the decomposition introduced in [7].

5.1 Random Clifford+T circuits

The first family of circuits we consider are Clifford+T circuits. We construct these by composing a given number (the T-count t) of *exponentiated Pauli unitaries*, of the form $\exp(-i(2k+1)\frac{\pi}{4}P)$ where P is a Pauli string, i.e. a tensor product of 1-qubit Pauli operators. The *weights* of the Pauli strings (the number of non-identity operators) is chosen randomly between 2 and 4, in order to mimic the structure common to quantum chemistry circuits, where the Hamiltonian has terms of weight 2–4.

We considered random 20-qubit Clifford+T circuits with T-counts varying from 1 to 43 (increasing in steps of 3). We calculated the amplitude of these circuits for a fixed input of $|+\rangle^{\otimes 20}$ and output of $\langle +|^{\otimes 20}$ (where $|+\rangle := \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$) using the two decomposition strategies. The running times shown in Figure 1 are averages over 10 runs for each T-count. We see already an order of magnitude improvement between the running time of the two approaches for a T-count of 43.

5.2 Random hidden-shift circuits

The second family of circuits we consider are hidden-shift circuits, whose description can be found in Refs. [4, 5]. These circuits are composed of H, Z, CZ and CCZ gates. The first three are readily translatable to ZX-diagrams, while the CCZ gate requires a decomposition that introduces T-spiders. The decomposition used here is one that uses 7 T-spiders for each CCZ gate.

We found the cat-state decompositions to be particularly efficient compared to the BSS decomposition for this class of circuits. We empirically found the BSS-based approach to be quite erratic, with running times varying between 0.005s and more than 1000s for circuits with the same T-count. This aspect seems to be dampened when using cat-state decompositions. To better gauge this phenomenon, we generated 125 20-qubit hidden-shift circuits with T-count 112, and simulated all the circuits using both approaches. The results, sorted by the simulation times for the cat-state method, can be seen in Figure 2. In the most extreme case, for simulating the same circuit, the cat-state-based decomposition ran in less than half a second, while the BSS-based one took more than 1 h 25 m.

The distribution of running times for both approaches can be found in Figure 3. As we are dealing with an inherently exponential process, it makes sense that perturbations have an exponential effect on the result. This is why we plotted the distributions on a logarithmic scale. In this same scale, it is possible to compute the variance of both distributions. We get a variance of $\sigma^2 \simeq 0.523$ for the cat-state approach, and a variance of $\sigma^2 \simeq 3.02$ for the BSS approach. This is strong evidence that the erratic behaviour is indeed dampened with our decomposition scheme, even when taking the exponential scaling into account.

This erratic behaviour being lifted with the cat-state strategy, we were able to run tests without the need for a time limit, which Ref. [15] set at 5 minutes. We re-ran the experiment that performed strong simulation of 100 random 50-qubit hidden-shift circuits with T-count 1400 using our new decompositions. We observed that 98% of the runs took less than 5 minutes, with the remaining circuits finishing within 6 minutes. This is to be compared with the BSS decomposition approach which in [15] obtained a 17% success rate (with the 5 minutes limit) even though it ran on a dedicated server with better CPU and parallelisation. The running time distribution is given in Figure 4.

6 Conclusion and Further Work

In this paper we built on the stabiliser rank quantum simulation method by combining it with a ZX-calculus simplification strategy and novel decompositions of entangled non-stabiliser states. We additionally introduced a new technique of *partial* magic state decompositions that only reduce T-count instead of eliminating all at once, which achieves the stabiliser rank upper bound of $2^{0.396t}$ found in [16] directly instead of just asymptotically. Our benchmarks shows that our technique is in the best case several orders of magnitude faster than the strategy of [15], which itself was already capable of simulating much larger circuits than the previous best. Additionally, our techniques seem to ‘smooth out’ the erratic runtimes of [15] allowing for more consistent simulation times.


For future work it would be interesting to investigate whether other classes of entangled states with an even smaller stabiliser rank could be found and used for this kind of exact simulation. It is also worth investigating how decompositions of these states can be used in methods for *approximate* simulations, like the stabiliser extent methods, or the norm estimation technique used in Ref. [4].

To sample from the random Clifford+T circuits we used the same technique as [15], which required a technique that doubles the T-count in the worst case. It would be interesting to see if the recent technique of calculating marginals without doubling [6] would help improve run times.


References

- 1 Scott Aaronson and Daniel Gottesman. Improved simulation of stabilizer circuits. *Physical Review A*, 70(5):052328, 2004. doi:10.1103/PhysRevA.70.052328.
- 2 Dorit Aharonov. A simple proof that Toffoli and Hadamard are quantum universal. *arXiv preprint quant-ph/0301040*, 2003.
- 3 Miriam Backens. The ZX-calculus is complete for stabilizer quantum mechanics. *New Journal of Physics*, 16(9):093021, 2014. doi:10.1088/1367-2630/16/9/093021.
- 4 Sergey Bravyi, Dan Browne, Pdraic Calpin, Earl Campbell, David Gosset, and Mark Howard. Simulation of quantum circuits by low-rank stabilizer decompositions. *Quantum*, 3:181, 9 2019. doi:10.22331/q-2019-09-02-181.
- 5 Sergey Bravyi and David Gosset. Improved classical simulation of quantum circuits dominated by Clifford gates. *Physical Review Letters*, 116(25), 6 2016. doi:10.1103/PhysRevLett.116.250501.
- 6 Sergey Bravyi, David Gosset, and Yinchun Liu. How to simulate quantum measurement without computing marginals. *arXiv preprint arXiv:2112.08499*, 2021.
- 7 Sergey Bravyi, Graeme Smith, and John A Smolin. Trading classical and quantum computational resources. *Physical Review X*, 6(2):021043, 2016. doi:10.1103/PhysRevX.6.021043.
- 8 Bob Coecke and Ross Duncan. Interacting quantum observables. In *Proceedings of the 37th International Colloquium on Automata, Languages and Programming (ICALP)*, Lecture Notes in Computer Science, 2008. doi:10.1007/978-3-540-70583-3_25.
- 9 Bob Coecke and Ross Duncan. Interacting quantum observables: categorical algebra and diagrammatics. *New Journal of Physics*, 13:043016, 2011. doi:10.1088/1367-2630/13/4/043016.
- 10 Alexander Cowtan, Silas Dilkes, Ross Duncan, Will Simmons, and Seyon Sivarajah. Phase Gadget Synthesis for Shallow Circuits. In Bob Coecke and Matthew Leifer, editors, *Proceedings 16th International Conference on Quantum Physics and Logic, Chapman University, Orange, CA, USA., 10-14 June 2019*, volume 318 of *Electronic Proceedings in Theoretical Computer Science*, pages 213–228. Open Publishing Association, 2020. doi:10.4204/EPTCS.318.13.
- 11 Ross Duncan, Aleks Kissinger, Simon Perdrix, and John van de Wetering. Graph-theoretic Simplification of Quantum Circuits with the ZX-calculus. *Quantum*, 4:279, 6 2020. doi:10.22331/q-2020-06-04-279.
- 12 Amar Hadzihasanovic, Kang Feng Ng, and Quanlong Wang. Two complete axiomatisations of pure-state qubit quantum computing. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '18*, pages 502–511, New York, NY, USA, 2018. ACM. doi:10.1145/3209108.3209128.
- 13 Emmanuel Jeandel, Simon Perdrix, and Renaud Vilmart. A Complete Axiomatisation of the ZX-calculus for Clifford+T Quantum Mechanics. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 559–568. ACM, 2018. doi:10.1145/3209108.3209131.
- 14 Aleks Kissinger and John van de Wetering. Reducing the number of non-Clifford gates in quantum circuits. *Physical Review A*, 102:022406, 8 2020. doi:10.1103/PhysRevA.102.022406.
- 15 Aleks Kissinger and John van de Wetering. Simulating quantum circuits with ZX-calculus reduced stabiliser decompositions, 2021. arXiv:2109.01076.
- 16 Hammam Qassim, Hakop Pashayan, and David Gosset. Improved upper bounds on the stabilizer rank of magic states. *Quantum*, 5:606, 12 2021. doi:10.22331/q-2021-12-20-606.
- 17 Yaoyun Shi. Both Toffoli and controlled-NOT need little help to do universal quantum computation. Preprint, 2002. arXiv:quant-ph/0205115.
- 18 John van de Wetering. ZX-calculus for the working quantum computer scientist. *arXiv preprint arXiv:2012.13966*, 2020.
- 19 Renaud Vilmart. A Near-Minimal Axiomatisation of ZX-Calculus for Pure Qubit Quantum Mechanics. In *2019 34th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–10, 2019. doi:10.1109/LICS.2019.8785765.


A Plots




■ **Figure 1** Runtime of random 20-qubit Clifford+T circuit simulations (avg of 10 runs per T-count).



■ **Figure 2** Logarithmic plot of the runtimes of simulating 125 different 20-qubit hidden-shift circuits with T-count 112, using the BSS decomposition and our strategy. Each datapoint corresponds to a circuit and they are sorted left-to-right according to simulation time using our approach.



■ **Figure 3** Distributions of running times (in seconds) of simulating 125 different 20-qubit hidden-shift circuits with T-count 112 using the BSS decomposition (left) and the cat-state decomposition (right).



■ **Figure 4** The time distribution of simulating 100 random 50-qubit hidden-shift circuits with T-count 1400 using our new decompositions.

