

Procedurally Optimised ZX-Diagram Cutting for Efficient T-Decomposition in Classical Simulation

Matthew Sutcliffe

Department of Computer Science
University of Oxford
Oxford, UK

matthew.sutcliffe@cs.ox.ac.uk

Aleks Kissinger

Department of Computer Science
University of Oxford
Oxford, UK

aleks.kissinger@cs.ox.ac.uk

A quantum circuit may be strongly classically simulated with the aid of ZX-calculus by decomposing its t T-gates into a sum of $2^{\alpha t}$ classically computable stabiliser terms. In this paper, we introduce a general procedure to find an optimal pattern of vertex cuts in a ZX-diagram to maximise its T-count reduction at the cost of the fewest cuts. Rather than reducing a Clifford+T diagram based on a fixed routine of decomposing its T-gates directly (as is the conventional approach), we focus instead on taking advantage of certain patterns and structures common to such circuits to, in effect, design by automatic procedure an arrangement of spider decompositions that is optimised for the particular circuit. In short, this works by assigning weights to vertices based on how many T-like gates they are blocking from fusing/cancelling and then appropriately propagating these weights through any neighbours which are then blocking weighted vertices from fusing, and so on. Ultimately, this then provides a set of weightings on relevant nodes, which can then each be cut, starting from the highest weighted down. While this is a heuristic approach, we show that, for circuits small enough to verify, this method achieves the most optimal set of cuts possible 71% of the time. Furthermore, there is no upper bound for the efficiency achieved by this method, allowing, in principle, an effective decomposition efficiency $\alpha \rightarrow 0$ for highly structured circuits. Even applied to random pseudo-structured circuits (produced from CNOTs, phase gates, and Toffolis), we record the number of stabiliser terms required to reduce all T-gates, via our method as compared to that of the more conventional T-decomposition approaches (namely [19], with $\alpha \approx 0.47$), and show consistent improvements of orders of magnitude, with an effective efficiency $0.1 \lesssim \alpha \lesssim 0.2$.

1 Introduction

Present-day quantum hardware is very limited, with few qubits and much noise [25]. Consequently, there are many classical techniques that are often employed to better optimise quantum circuits and/or to verify the behaviour of quantum hardware and software. A particularly useful tool for facilitating this is ZX-calculus [11, 30, 1, 12, 20], which allows quantum circuits to be expressed and simplified graphically with the use of known rewriting rules. This has been utilised for various problems in the field, including notably optimisation [14, 13, 2, 5, 15, 23, 24] and classical simulation [19, 21, 31, 9, 10, 8, 22]. On the latter problem, for instance, wherein one wishes to compute the probabilities of particular measurement outcomes of a quantum circuit, ZX-calculus can be utilised to re-express a large ‘Clifford+T’ circuit (which is notoriously inefficient to simulate classically) as a sum of ‘Clifford’ circuits (which *are* efficient to classically simulate) [19].

Re-expressing a Clifford+T circuit as sum of Clifford circuits in this way relies upon known decompositions of sets of costly ‘T-gates’ into cheap ‘Clifford’ terms. There are many such decompositions that have been discovered [7, 6, 19, 21], of varying efficiencies, and the conventional strategies [19], perhaps unsurprisingly, tend to opt for the most efficient (i.e. those which translate a set of T-gates into

the fewest number of Clifford terms). However, the apparent efficiencies of these decompositions can be misleading. In fact, it has been noted that applying apparently *less* efficient decompositions in certain circumstances can result in *more* efficient results overall (that is, fewer Clifford terms in exchange for removing all T-gates of a circuit) [9]. Knowing when this can be applicable on local scales is often fairly straightforward. However, what is seldom considered is how appropriate applications of these less efficient decompositions can be found when looking on broader scales.

Specifically, in this paper, we demonstrate how drastic improvements to the overall efficiency can be attained by applying such decompositions to inconspicuous gates in a circuit with no obvious or immediate local benefit for doing so. We then present a means by which one can procedurally analyse a given circuit to determine where these optimal gates for decomposition are. We do this based on a heuristic of weighing vertices due to their immediate local benefit of applying such a decomposition, and then propagating these weights through their neighbours as appropriate. Lastly, we showcase the effectiveness of our method by comparing its overall efficiency at fully decomposing various Clifford+T circuits, of various numbers of T-gates, against the more traditional methods [19].

2 Background

2.1 ZX-calculus

A very useful notation with which to express quantum circuits is that provided by *ZX-calculus* [11, 30, 1], wherein all operations are expressed as *spiders* (phase rotations) about either the Z-axis or X-axis of the Bloch sphere, connected via *edges* (or *wires*). Circuits expressed in this notation are known as *ZX-diagrams* and, by convention, Z-spiders are green and X-spiders red - in either case with the angle of rotation written within (or left blank if zero), as follows:

$$\begin{aligned}
 \text{Z-spider}(\alpha) &:= |0\dots 0\rangle\langle 0\dots 0| + e^{i\alpha}|1\dots 1\rangle\langle 1\dots 1| \\
 \text{X-spider}(\alpha) &:= |+\dots+\rangle\langle +\dots+| + e^{i\alpha}|-\dots-\rangle\langle -\dots-|
 \end{aligned}$$

In addition to these two types of spiders, for convenience the Hadamard gate is often also included explicitly, as a yellow box, though this too may be decomposed (via the Euler decomposition) into spiders as such:

$$\text{---} \square \text{---} \equiv e^{-i\frac{\pi}{4}} \text{---} \text{O}(\frac{\pi}{2}) \text{---} \text{O}(\frac{\pi}{2}) \text{---} \text{O}(\frac{\pi}{2}) \text{---} \equiv \text{---} \text{---} \text{---}$$

As shown here, an edge containing a Hadamard (referred to, unsurprisingly, as a *Hadamard edge*) may alternatively be expressed as a dashed blue line.

From these basic components, any quantum circuit may be expressed as a ZX-diagram. Indeed, completeness for the Clifford gateset is achieved provided phases of $\frac{n\pi}{2}$ are allowed (where $n \in \mathbb{Z}$). (Note also that, as phases correspond to rotations, they are all of modulo 2π .) To extend the completeness to the Clifford+T gateset (including T-gates and Toffoli gates), the resolution must be expanded to support phases of $\frac{n\pi}{4}$.

2.2 ZX-diagram rewriting

The major benefit of ZX-calculus is that it contains a set of well-defined *rewriting rules* which outline how certain structures within ZX-diagrams may be equivalently written in simpler forms - allowing a circuit to be simplified while maintaining its behaviour. In particular, the fundamental rewrite rules are as outlined in figure 1.




Figure 1: A set of the basic rewriting rules of ZX-calculus [19], where Greek letters denote arbitrary real variables, $[0, 2\pi)$, and Latin letters denote arbitrary boolean variables, $\{0, 1\}$. Note that the rules still apply if all the spider colours are inverted. These rules are known, in column-major order, as (a) spider fusion, (b) colour change, (c) π -commutation, (d) bialgebra, (e) state copy (where n is the number of output edges), (f) identity removal, (g) Hadamard cancellation.

From these basic rules, one may derive a number of more complex rules. In particular, *local complementation* and *pivoting*, outlined in figure 2, prove to be extremely useful in reducing ZX-diagrams. A more thorough collection of derived rules may be found in [30].

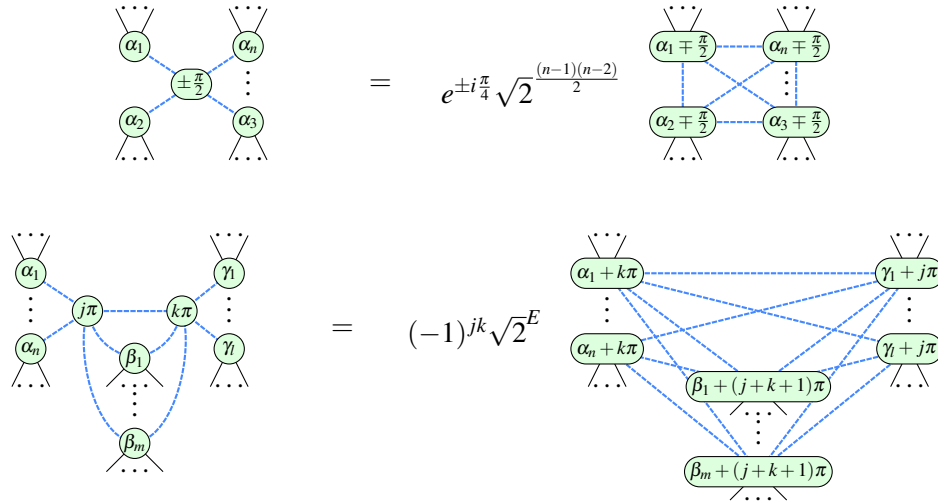


Figure 2: Two important derived ZX-calculus rewrite rules [19], namely (top) local complementation and (bottom) pivoting. Here, $E = (n - 1)m + (l - 1)m + (n - 1)(l - 1)$.

A particular subset of ZX-diagrams are those with no open input or output wires. These are known as *scalar diagrams*, and the rules above are sufficient to reduce such diagrams, of the Clifford gateset, to

complex scalars, making use of the following scalar relations:

$$\textcircled{\alpha} = 1 + e^{i\alpha} \quad \textcircled{\alpha} \textcircled{a\pi} = \sqrt{2}e^{ia\alpha}$$

2.3 Classical simulation

To verify the behaviour of complex quantum algorithms, especially when present-day quantum hardware is insufficient to run them, one may turn to classical simulation. There are two classes of this - namely *strong* simulation and *weak* simulation. For the purposes in this paper, we will focus on the former, wherein the aim is to determine the probability of a specific measurement outcome. In short, this may be done with ZX-calculus by ‘plugging’ the inputs and outputs with spiders corresponding to the desired states and measurements to produce a *scalar* ZX-diagram, which may then be fully reduced to a simple complex scalar value. This result will represent its amplitude, A , and thus relates to its probability, $|A|^2$.

Notably, Clifford diagrams (those restricted to phases of $\frac{n\pi}{2}$) are very efficient to simulate classically. This manifests itself in ZX-calculus in that such diagrams, given no open inputs or outputs, may be fully reduced to a scalar via the rewriting rules highlighted above. Clifford+T diagrams, on the other hand, are notoriously inefficient to simulate classically. While the rewriting rules may remove a number of T-gates from a ZX-diagram, it is generally unable to remove them all, meaning such diagrams are unable to fully reduce to a scalar. To deduce the scalar amplitude in such cases then, one must make use of decompositions to translate the Clifford+T diagram into a sum of efficiently-reducible Clifford diagrams. The inefficiency lies in the fact that the number of such summand Clifford terms one attains in place of a Clifford+T diagram scales exponentially with the number of T-gates. This is typically quantified with a parameter α which represents the efficiency of the precise decomposition(s) used, given it translates a Clifford+T diagram of t T-gates into a sum of $2^{\alpha t}$ Clifford terms. So, smaller values of α describe more efficient decompositions.

For instance, the decomposition presented by Bravyi, Smith, and Smolin [7], and expressed in ZX-calculus terms by Kissinger and van de Wetering [19], allows sets of 6 T-gates to be replaced with a sum of 7 Clifford terms, according to figure 3. As such, this ‘‘BSS’’ decomposition scales as $7^{t/6} \approx 2^{0.468t}$, hence $\alpha \approx 0.468$. The current state of the art T-decompositions, meanwhile, achieve $\alpha \approx 0.396$ [26].




Figure 3: The BSS decomposition [7], relating a set of 6 T-gates to a sum of 7 Clifford terms, expressed in ZX-diagrams as per [19].

As outlined in [19], after every decomposition, some further reduction in T-count may be facilitated by Clifford simplification via the rewriting rules, such that the $2^{\alpha t}$ actually represents an upper-bound estimate of the number of terms ultimately produced. Moreover, in their work they essentially select

groups of 6 T-gates for decomposition arbitrarily, without regard for whether certain groupings may be more likely to allow further simplification after decomposition. Lastly, while their method relies primarily on the BSS decomposition, when the number of T-gates remaining in any graph falls below 6, they fall back on other known decompositions.

3 Method

3.1 Graph cutting

While Clifford+T diagrams are typically reduced via such T-gate decompositions as that of Bravyi, Smith, and Smolin, highlighted above, paying close attention to the structures inherent in a given ZX-diagram can reveal that blindly decomposing T-gates in an essentially random order is not necessarily optimal. In fact, decomposing conveniently positioned *Clifford* spiders - even with an apparently less efficient decomposition - can actually produce more efficient results.

From the definitions of Z- and X- spiders in section 2.1, one can infer the basic relation:

$$\begin{array}{c} \vdots \\ \alpha \\ \vdots \end{array} \approx \begin{array}{c} \vdots \\ \text{red dots} \\ \vdots \end{array} + e^{i\alpha} \begin{array}{c} \pi \quad \pi \\ \vdots \\ \pi \quad \pi \end{array}$$

Note that the equality here is up to some global scalar factor (neglected for brevity). This acts as a very simple decomposition for any arbitrary spider. Herein, the act of applying this particular decomposition will be referred to as ‘*cutting*’ the graph, as structurally it behaves like physically slicing a vertex from its edges. Superficially, this may not seem particularly useful as, applied to T-gates, this doubles the number of terms for each one (hence has a very poor efficiency of $\alpha = 1$). However, by cutting appropriate spiders (not necessarily even T-like spiders), this can allow some T-gates to cancel, along both paths of the decomposition.

Specifically, consider a subgraph consisting of two T-like spiders sandwiched between a CNOT. Cutting one end of the CNOT (the side opposite the T-gates) will produce two terms, like so:

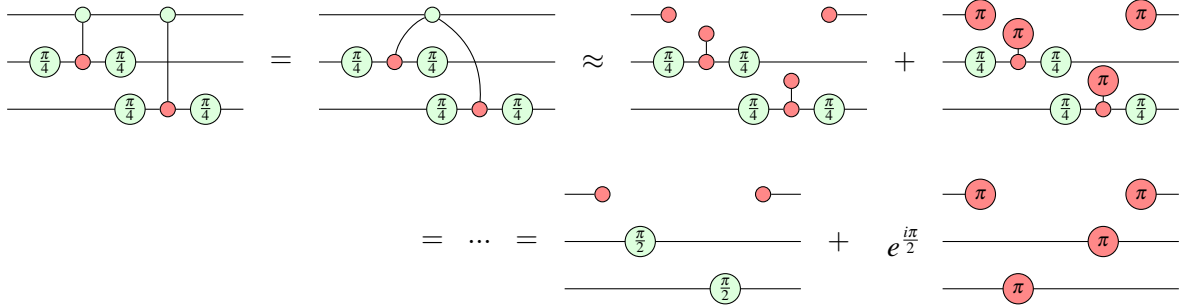
An observant reader may notice that both of these terms, after a little simplification via the rewriting rules, may fuse their T-gates into a Clifford. The left-hand term, for instance, simplifies as follows:

The right-hand term simplifies similarly, except utilising π -commutation in place of identity removal:

Consequently, via a cut of a Clifford spider the T-count has been reduced by 2. This gives a decomposition efficiency of $\alpha = 0.5$. Moreover, if the green end of the CNOT had itself taken a T-like phase, then that too would have been removed (specifically, converted into a scalar factor directly by the decomposition). In that case, one would have observed a better efficiency of $\alpha = 1/3$.

3.2 CNOT-grouping

One may further recognise that it is possible for any number of these T-CNOT-T arrangements to be aligned such that their respective CNOTs may fuse. In such cases, each pair of T-like spiders may fuse to a Clifford as above, though still just requiring one vertex cut overall. As a simple example, the following shows how two sets might be grouped and reduced in this way:




Here, the T-count has been reduced by 4, still at the expense of just one vertex cut (hence 2 Clifford terms). This corresponds to $\alpha = 0.25$. It should be clear from this point that this reasoning could be extended ad infinitum, with arbitrarily many T-CNOT-T subgraphs being fusible and hence arbitrarily many T-gates being reducible to Clifford at the expense of just one cut, giving $\alpha \rightarrow 0$.

A similar concept to this was recognised in [9] (albeit notated in a different fashion), wherein, for a single decomposition utilising π -commutations, T-count reductions of up to 286 were found on SAT counting [4, 3] ZX-diagrams (giving $\alpha \approx 0.0035$). However, this work utilised a very simple heuristic amounting to prioritising decompositions of T-gates with the maximum number of immediately relevant connections. Perhaps surprisingly, however, this naïve heuristic seldom produces the most efficient solutions and indeed can often be extremely suboptimal. Moreover, their approach to recognising instances where these cuts are applicable is essentially limited to those where the decomposition is to be applied to a T-like spider, rather than any arbitrarily-phased spider as per the means outlined above. And lastly, by acting on ZX-diagrams *after* they have undergone full Clifford simplification, their approach risks losing much of the graph's structure and thus missing relevant patterns.

3.3 Cutting in tiered structures

Just as a CNOT may be the only thing standing in the way of two or more T-like spiders from fusing, so too may a CNOT be the only obstacle preventing some T-CNOT-T sandwiches from grouping. Consider, as a prime example, the structured circuit that follows, where the Z-spiders here have been labelled for easy reference:



Immediately, one might recognise 7 T-CNOT-T sandwiches, centred on vertices #2, #4, #6, #9, #11, #13, and #15, respectively. Notably, however, there is some clashing here, in that - for example - the

sandwich centred on #2 and that centred on #4 are mutually exclusive, as they share one of their T-gates. Indeed, there are a number of possible ways of reducing this circuit via the means outlined so far in this paper.

For example, cutting vertices #4, #9, and #13 would allow 6 T-gates to reduce to Clifford. Alternatively, cutting #2, #6, #11, and #15 would reduce all 8 T-gates, at the cost of 4 cuts (hence $\alpha = 0.5$). But, the best solution here would be to firstly cut vertex #8 - even though this doesn't immediately allow any T-gates to reduce - as this then allows vertices #2, #6, #11, and #15 to fuse into one. Cutting this newly fused vertex then would allow all 8 T-gates to reduce. So, this solution would reduce the T-count entirely (all 8 T-gates) at the cost of just 2 cuts (hence $\alpha = 0.25$). And of course, this reasoning could be extended for higher tiers, where the optimal initial cuts are two, three, or more, steps away from any T-gates.

Similarly, there may be instances where *multiple* CNOTs are directly blocking some set of T-gates from fusing. In such cases, the likelihood that the cutting all of those CNOTs would be worthwhile to reduce the T-gates they block will be determined by the ratio of the CNOTs to T-gates involved, as well as whether some of those CNOTs are considered worthwhile cuts in their own right, with regard to any other T-gates that they alone may be blocking. We can call such groups of spiders 'spousal', with respect to the children spiders they are collectively blocking from reducing.

Evidently, therefore, selecting which vertices to cut, and indeed in which order, is a very intricate task (though the former is more important as suboptimality in the latter can be corrected for via a slight modification and parametric analysis, as detailed in appendix A). Naïvely tackling this problem via an exhaustive, brute force, approach would require checking the reduction achieved by every possible combination of vertex cuts. This is obviously infeasible for large-scale graphs, as the time complexity scales exponentially with the number of vertices. Consequently, a heuristic approach is desired. On this note, as we have shown, simply prioritising vertices which are directly blocking the most number of T-like pairs from fusing is not generally optimal. Rather, it is preferential to look at the whole picture and determine the optimal cuts on higher tiers.

3.4 Optimised cutting procedure

The solution we propose is a procedure based on assigning weights to vertices, determined by how many T-like gates they are preventing from fusing to Clifford, and then propagating these weights through any neighbours which are then preventing weighted vertices from fusing, and so on up the tiers. Particular care is given to balance the weightings appropriately, especially in places where multiple 'spousal' cuts are required to facilitate a fusion of their children.

Note that we may label the weight of a vertex (i.e. spider) v , for a particular tier t , as w_v^t , such that, for instance, w_{12}^2 refers to the weight of vertex #12 with respect to tier #2. Given this, the procedure steps are as follows:

1. Partially simplify the circuit such that any instances of spider fusion are applied (so no like-coloured spiders remain directly connected via a solid edge) and any π -phase spiders are pushed to one side or into CNOTs via the π -commutation and fusion rules. Then, let $t = 0$ and assign an initial weight of 0 to every spider (i.e. let $w_v^0 = 0 \forall v$).
2. For any pair of T-like spiders that are separated by k (that is, one or more) CNOTs, add $2/k$ to the weights of the opposite ends of each of those CNOTs. (For any given CNOT, take care not to count a particular T-spider more than once.) Now, for instance, any CNOT that is preventing 2 T-spiders from fusing to Clifford will have a corresponding weight of 2, and any CNOT preventing,

collectively, 4 T-spiders from reducing to Cliffords will have a corresponding weighting of 4, etc. Similarly, if, for instance, 3 CNOTs are collectively blocking a single pair of T-spiders from fusing, then each will have a weighting of $2/3$.

3. Increment $t \leftarrow t + 1$. Then, similar to step 2, for any *weighted* vertex v of the previous tier (i.e. any v for which $w_v^{t-1} \geq 0$) that is separated from fusing with another weighted vertex of *any* lower tier (i.e. any v for which $w_v^u \geq 0$ for any $u < t$) by k (that is, one or more) CNOTs, add $\gamma(w_v^{t-1})/k$ to the weight of the opposite ends of each of those CNOTs. Here, the γ function normalises a given weighting to the range $[0, 1]$, such that (crudely speaking) 0 roughly means “very unlikely to be a worthwhile cut” and 1 “very likely to be a worthwhile cut”: $\gamma(w) := \min(\frac{w}{2}, 1)$.
4. Repeat step 3 until no new changes are made (that is, until no weightings, w_v^t , are found for any v , given the current t). At this point, one will be left with weightings for every vertex, for every tier ($w_v^t, \forall v, t$). From this, one can trivially extract the relevant data, namely, for every vertex v , its *maximum* weight w_v^t (for any t) and, respectively, the maximum t for which the vertex has a weight (i.e. largest t for which $w_v^t \geq 0$). One may label the maximum weight of a vertex, W_v , and its maximum tier for which it has a non-zero weight, T_v .
5. Among the vertices with the largest recorded T , namely T_{max} , select the one with the greatest max weight W_v (i.e. select vertex V s.t. $W_V \geq W_v \forall v$ s.t. $T_v = T_V = T_{max}$). (Note that for this step only, we may temporarily add 1 to the weight of any vertex of a T-like phase.) If this weight is below 2 (hence $\gamma(W_V) \leq 1$, implying the cut would not likely be worthwhile), then search instead among the vertices of the lower tier (i.e. for which $T_v = T_{max-1}$), until an appropriate vertex is found for which $W_V \geq 2$. This is the vertex which the heuristic has concluded is likely an optimal choice to cut. As such, cut this vertex (i.e. decompose it into two branches as per section 3.1).
6. Partially simplify each branch, without compromising the graph structure. Specifically, push any new π -phase spiders to one side, and/or into CNOTs, via repeated applications of the π -commutation rule (and fusion), and thereafter apply the fusion rule until no like-coloured spiders remain connected via a solid edge. Moreover, when fusing weighted spiders, update their combined weight accordingly as the sum of their respective max weights (i.e. in fusing vertex B into vertex A , the former is removed along with its recorded weightings and tier data and the latter is updated as $W_A \leftarrow W_A + W_B$). Similarly, the max tier of the newly fused vertex takes that of the larger of the two fused vertices (i.e. $T_A \leftarrow \max(T_A, T_B)$). Moreover, recalculate the weightings on any vertices whose children have been altered by this partial simplification.
7. Repeat steps 5 and 6 until no further cuts are made (or the number of T-spiders ≤ 2). If any T-like spiders remain, then these may be decomposed via a typical T-decomposition (e.g. the BSS decomposition outlined in section 2.3).

Applying this procedure to a Clifford+T ZX-diagram will result in a set of Clifford ZX-diagrams whose sum equals the original. For scalar diagrams (those with no open input or output wires), each summand will simply be a scalar term, and thus the overall sum will be likewise [19].

Python code, based on the PyZX package [16, 17], that implements this procedure may be found at <https://github.com/mjsutcliffe99/ProcOptCut> [28], and a step-by-step illustrative example of this procedure in action is shown therein.

4 Results

4.1 Complexity and efficiency

The exact runtime complexity of the procedure is difficult to discern as it depends on the density of the graph (i.e. typical number of neighbours the vertices have) and its number of ‘tiers’, T_{max} , (which is heuristically determined). Very reasonably, one could assume $T_{max} \ll V$ (for a non-trivial number of vertices, V), and so a crude upper-bound (taking, very unrealistically, maximal density) may be given by $O(V^2)$. Even this gross over-estimate of the runtime complexity shows that, compared to the number of stabiliser terms produced ($2^{\alpha t}$ for t T-gates and some $\alpha < 1$), the procedure’s runtime is negligible, so that any improvements it offers to the number of the stabiliser terms can be taken as is.

To test the effectiveness of the proposed method, we explored how often - for circuits small enough to verify - it was able to find the *most* optimal set of vertex cuts, rather than simply *an* optimal set. Specifically, we considered many random, non-trivially structured (see appendix B) ZX-diagrams of 16 internal Z-spiders or fewer and measured the number of stabiliser terms required to remove all T-spiders via the method outlined in section 3.4. In each case, we also tested every possible combination of vertex cuts on the internal Z-spiders (applying the reasoning of appendix A to correct for any suboptimal cut ordering). Thus, we could determine with certainty the most optimal set of vertex cuts (on Z-spiders) and correspondingly the number of stabiliser terms produced (or, by extension, the effective α) and compare this, in each case, to the result achieved by our method. In this way, we observed that our method found the most optimal set of cuts possible (on Z-spiders) 71% of the time. Even so, in every case in which the method failed to find the most optimal solution, it still invariably found a very good solution with an α usually only marginally above that of the best. It is worth noting, however, that we cannot be certain that this would translate to a similar success rate for much larger circuits (i.e. those too big to verify by brute force), where greater levels of structure (and indeed chaos) are possible. As such, to extend our analysis we compared our method also to that of the more typical approach of T-decompositions a la [19], of which we outline our results ahead.

4.2 Experimental measurements for random circuits

Applying the procedure outlined in section 3.4 to a Clifford+T ZX-diagram will produce an exact result (e.g. for computing the probability amplitude of a Clifford+T diagram for classical simulation). However, in benchmarking the efficiency of the method, one is more interested in the *number* of stabiliser terms the original diagram is decomposed to, rather than the precise value of the sum of these terms. With this in mind, one can run a ‘*blind*’ version of the procedure to compute an (upper-bound) estimate of the number of resulting terms in linear time. As well as not providing the final numerical result, this blind version of the procedure also misses out on inter-step ZX-calculus simplifications, and hence typically overestimates the number of resulting terms (and thus is an upper-bound estimate). The appropriate modifications for this version of the procedure are outlined and justified in appendix C.

Similarly, one can likewise calculate an upper-bound estimate of the number of terms produced after fully decomposing a Clifford+T diagram via the BSS-based approach of [19]. This is much more simple to do as the BSS decomposition efficiency is known in advance, namely $\alpha \approx 0.468$. Thus, an estimate of the number of resulting terms this method achieves for a given diagram of T-count t is given by $2^{0.468t}$, which can be computed trivially in constant time. As before, this does not account for the inter-step ZX-calculus simplifications and hence is an approximate upper-bound estimate.

For various randomly generated ZX-diagrams (constructed from CNOTs, $\frac{n\pi}{4}$ phase gates, and Tof-

folis as per appendix B), we measured the number of stabiliser terms required to remove all T-spiders via the approach of Kissinger and van de Wetering [19] versus the method we propose in section 3.4. For computationally feasible T-counts (those computable within a few minutes), we can compute these exactly with both methods, and for larger (computationally infeasible) T-counts, we instead rely on the upper-bound estimations outlined above. For each of these measurements, we can also infer the overall effective decomposition efficiency, α , by simply taking $\alpha = \frac{1}{t} \log_2 n$, where t is the initial T-count of the circuit (after full Clifford simplification) and n is the number of stabiliser terms to which it is reduced. The results are illustrated in figure 4. (Note that the initial T-counts are all taken as that achieved after full Clifford simplification.)




Figure 4: (a) The number of stabiliser terms, n , produced after decomposing all t T-gates, and correspondingly (b) the effective overall decomposition efficiency α (given by $\alpha = \frac{1}{t} \log_2 n$) for numerous randomly generated (see appendix B) pseudo-structured Clifford+T circuits of various T-counts. In each case, we measure experimentally the exact results, as well as approximate upper-bound estimations, achieved by both the conventional (“BSS”) method of [19] and the procedural approach (“proc”) presented in this paper.

We observe that our procedural method, applied to circuits of this type, is very effective at minimising the number of resulting stabiliser terms, compared to relying predominantly on naïve applications of the BSS decomposition. The exact effectiveness of our method can vary quite substantially from circuit to circuit, depending on how much structure they embed. The direct BSS approach, meanwhile, tends to be a lot more consistent, though can also vary somewhat depending on how much inter-step simplification it is able to undertake. Regardless, the procedural method very consistently shows magnitudes of improvement on these circuits, allowing above even double the T-counts to be computed within the same time frame.

This is reflected too in the effective α measurements, where, for circuits of this type, our method offers typically $0.1 \lesssim \alpha \lesssim 0.2$ for such pseudo-structured circuits. The results are less remarkable on trivially small T-counts as such small circuits have little room for much structure of which to take advantage. Evidently, therefore, this method is, as intended, very well optimised for decomposing Clifford+T circuits that exhibit some structure, enabling such circuits of much larger T-counts to be computed within reasonable time frames.

For full context, it is worth remembering two points when considering these results. Firstly, the random circuits generated for these experiments were those with some inherent localised structural elements. And secondly, while the runtime of the relevant steps of the procedure, after each decomposition,

is very quick, it does nevertheless contribute time to each term that is not present in the BSS method, so that the difference in numbers of terms produced by the two methods does not translate 1-to-1 to a difference in runtime. In other words, computing n terms via the procedural method may be a little slower than computing n terms via the BSS method (although this difference will not be more than some small factor, so that the improvement offered by the reduced number of terms still vastly outweighs this offset).

All of the above experiments may be repeated from the corresponding Jupyter notebook, hosted on Github [28].

5 Conclusion

We demonstrated how a procedure could be designed to analyse the structure of any given Clifford+T quantum circuit in order to determine an optimised set of vertex cuts to efficiently decompose it to a sum of (classically computable) Clifford terms. This is contrary to the more conventional approach [19], which applies decompositions essentially arbitrarily, without such regard to the specific structures involved.

Specifically, we show that our method is very effective at finding optimal sets of vertex cuts, with a 71% success rate at finding the *most* optimal set when applied to random small semi-structured circuits. We further show that our method is very efficient at decomposing even larger such circuits, consistently outperforming the more conventional approach [19] by orders of magnitude, with regard to the number of resulting terms (and by extension the runtime). What this means in practise is that our method could allow for classical simulation, within a reasonable time frame, of Clifford+T circuits of more than double the T-count as could be achieved with the conventional methods. This is hugely relevant to verifying the behaviour of quantum software and hardware for present-day NISQ (noisy intermediate-scale quantum) devices [25].

While already very effective, there are many ways in which the method outlined in this paper could be improved - many stemming from the rigid scope with which it applies and propagates weights. After all, T-gate fusion facilitated by cutting a CNOT is just one way in which the T-count of a circuit may be reduced. It would be worth considering also the potential of cuts to remove T-gates by pushing them into the scalar factor (e.g. via the state copy rule) or to partition small segments from the graph. The method ought also to consider simplification and weight propagation laterally, rather than solely through the lens of ‘tiers’. Indeed, the ‘partial simplification’ strategy we utilise (to simplify while maintaining the structure) does not take into account some vital steps in the normal ‘full reduce’ [16] function (namely pivoting - i.e. on CNOTs - and local complementation). We suspect this is largely the reason we didn’t observe an even higher success rate in verifying how often the method was able to find the best solutions on small circuits. Moreover, a more robust analysis could determine appropriate weightings *without* the need for the circuit to be expressed in a very rigid graph-like form, such that further simplification between steps could be enabled and, for instance, Toffolis could be expressed as phase gadgets (so that we are not restricted due to the arbitrary choice of which way around to decompose each Toffoli’s control qubits). And lastly, of course, one could consider cuts on X-spiders as well as just Z-spiders (this might be particularly relevant if there are many Hadamards involved, resulting in many X-spiders of T-like phase).

There are also a number of ways in which this concept, more broadly, could be improved, such as developing newer and better heuristics - perhaps even different heuristics for different types of circuit (e.g. dense circuits, or those with many Toffolis, etc.). Nevertheless, we demonstrate very clearly how analysing the circuit structure and applying decompositions discriminately can offer vastly more efficient

results than simply decomposing the T-spiders directly with a decomposition that has a better *immediate* efficiency.

References

- [1] Miriam Backens, Simon Perdrix & Quanlong Wang (2017): *A Simplified Stabilizer ZX-calculus*. *Electronic Proceedings in Theoretical Computer Science* 236, p. 1–20, doi:10.4204/eptcs.236.1.
- [2] Niel de Beaudrap, Xiaoning Bian & Quanlong Wang (2020): *Fast and Effective Techniques for T-Count Reduction via Spider Nest Identities*. In Steven T. Flammia, editor: *15th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2020)*, *Leibniz International Proceedings in Informatics (LIPIcs)* 158, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, pp. 11:1–11:23, doi:10.4230/LIPIcs.TQC.2020.11.
- [3] Niel de Beaudrap, Aleks Kissinger & Konstantinos Meichanetzidis (2021): *Tensor Network Rewriting Strategies for Satisfiability and Counting*. *Electronic Proceedings in Theoretical Computer Science* 340, p. 46–59, doi:10.4204/eptcs.340.3.
- [4] Lucas Berent, Lukas Burgholzer & Robert Wille (2022): *Towards a SAT Encoding for Quantum Circuits: A Journey From Classical Circuits to Clifford Circuits and Beyond*. doi:10.48550/arXiv.2203.00698.
- [5] Agustín Borgna, Simon Perdrix & Benoît Valiron (2021): *Hybrid quantum-classical circuit simplification with the ZX-calculus*. In Hakjoo Oh, editor: *Programming Languages and Systems*, Springer International Publishing, Cham, pp. 121–139, doi:10.1007/978-3-030-89051-3_8.
- [6] Sergey Bravyi, Dan Browne, Pádraic Calpin, Earl Campbell, David Gosset & Mark Howard (2019): *Simulation of quantum circuits by low-rank stabilizer decompositions*. *Quantum* 3, p. 181, doi:10.22331/q-2019-09-02-181.
- [7] Sergey Bravyi, Graeme Smith & John A. Smolin (2016): *Trading classical and quantum computational resources*. *Physical Review X* 6(2), p. 021043, doi:10.1103/PhysRevX.6.021043.
- [8] Tristan Cam & Simon Martiel (2023): *Speeding up quantum circuits simulation using ZX-Calculus*. *arXiv preprint arXiv:2305.02669*.
- [9] Julien Coudi (2022): *Cutting-Edge Graphical Stabiliser Decompositions for Classical Simulation of Quantum Circuits*. Master’s thesis, University of Oxford. Available at <https://www.cs.ox.ac.uk/people/aleks.kissinger/theses/coudi-thesis.pdf>.
- [10] Julien Coudi & John van de Wetering (2022): *Classically Simulating Quantum Supremacy IQP Circuits through a Random Graph Approach*. *arXiv preprint arXiv:2212.08609*.
- [11] Bob Coecke & Ross Duncan (2011): *Interacting quantum observables: categorical algebra and diagrammatics*. *New Journal of Physics* 13, p. 043016, doi:10.1088/1367-2630/13/4/043016.
- [12] Bob Coecke & Aleks Kissinger (2017): *Picturing Quantum Processes*. Cambridge University Press, doi:10.1017/9781316219317.
- [13] Alexander Cowtan, Silas Dilkes, Ross Duncan, Will Simmons & Seyon Sivarajah (2020): *Phase Gadget Synthesis for Shallow Circuits*. In Bob Coecke & Matthew Leifer, editors: *Proceedings 16th International Conference on Quantum Physics and Logic, Chapman University, Orange, CA, USA., 10-14 June 2019*, *Electronic Proceedings in Theoretical Computer Science* 318, Open Publishing Association, pp. 213–228, doi:10.4204/EPTCS.318.13.
- [14] Ross Duncan, Aleks Kissinger, Simon Perdrix & John van de Wetering (2020): *Graph-theoretic Simplification of Quantum Circuits with the ZX-calculus*. *Quantum* 4, p. 279, doi:10.22331/q-2020-06-04-279.
- [15] Stefano Gogioso & Richie Yeung (2023): *Annealing Optimisation of Mixed ZX Phase Circuits*. In Stefano Gogioso & Matty Hoban, editors: *Proceedings 19th International Conference on Quantum Physics and Logic, Wolfson College, Oxford, UK, 27 June - 1 July 2022*, *Electronic Proceedings in Theoretical Computer Science* 394, Open Publishing Association, pp. 415–431, doi:10.4204/EPTCS.394.20.

- [16] Aleks Kissinger & John van de Wetering: *PyZX*. Available at <https://github.com/Quantomatic/pyzx>.
- [17] Aleks Kissinger & John van de Wetering (2020): *PyZX: Large Scale Automated Diagrammatic Reasoning*. In Bob Coecke & Matthew Leifer, editors: Proceedings 16th International Conference on *Quantum Physics and Logic*, Chapman University, Orange, CA, USA., 10-14 June 2019, *Electronic Proceedings in Theoretical Computer Science* 318, Open Publishing Association, pp. 229–241, doi:10.4204/EPTCS.318.14.
- [18] Aleks Kissinger & John van de Wetering (2020): *Reducing the number of non-Clifford gates in quantum circuits*. *Physical Review A* 102(2), doi:10.1103/physreva.102.022406.
- [19] Aleks Kissinger & John van de Wetering (2022): *Simulating quantum circuits with ZX-calculus reduced stabiliser decompositions*. *Quantum Science and Technology* 7(4), p. 044001, doi:10.1088/2058-9565/ac5d20.
- [20] Aleks Kissinger & John van de Wetering (2023): *Picturing Quantum Software [Preprint]*.
- [21] Aleks Kissinger, John van de Wetering & Renaud Vilmart (2022): *Classical Simulation of Quantum Circuits with Partial and Graphical Stabiliser Decompositions*. In François Le Gall & Tomoyuki Morimae, editors: *17th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2022)*, *Leibniz International Proceedings in Informatics (LIPIcs)* 232, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, pp. 5:1–5:13, doi:10.4230/LIPIcs.TQC.2022.5.
- [22] Mark Koch, Richie Yeung & Quanlong Wang (2023): *Speedy Contraction of ZX Diagrams with Triangles via Stabiliser Decompositions*. *arXiv preprint arXiv:2307.01803*.
- [23] Tommy McElvanney & Miriam Backens (2023): *Flow-preserving ZX-calculus Rewrite Rules for Optimisation and Obfuscation*. In Shane Mansfield, Benoit Valiron & Vladimir Zamdzhiev, editors: *Proceedings of the Twentieth International Conference on Quantum Physics and Logic, Paris, France, 17-21st July 2023*, *Electronic Proceedings in Theoretical Computer Science* 384, Open Publishing Association, pp. 203–219, doi:10.4204/EPTCS.384.12.
- [24] Maximilian Nägele & Florian Marquardt (2023): *Optimizing ZX-Diagrams with Deep Reinforcement Learning*. *arXiv preprint arXiv:2311.18588*.
- [25] John Preskill (2018): *Quantum Computing in the NISQ era and beyond*. *Quantum* 2, p. 79, doi:10.22331/q-2018-08-06-79.
- [26] Hammam Qassim, Hakop Pashayan & David Gosset (2021): *Improved upper bounds on the stabilizer rank of magic states*. *Quantum* 5, p. 606, doi:10.22331/q-2021-12-20-606.
- [27] Matthew Sutcliffe: *ParamZX*. Available at <https://github.com/mjsutcliffe99/ParamZX>.
- [28] Matthew Sutcliffe: *ProcOptCut*. Available at <https://github.com/mjsutcliffe99/ProcOptCut>.
- [29] Matthew Sutcliffe & Aleks Kissinger (2024): *Fast classical simulation of quantum circuits via parametric rewriting in the ZX-calculus [Preprint]*.
- [30] John van de Wetering (2020): *ZX-calculus for the working quantum computer scientist*. doi:10.48550/arXiv.2012.13966.
- [31] Robert Wille, Lukas Burgholzer, Stefan Hillmich, Thomas Grurl, Alexander Ploier & Tom Peham (2022): *The Basis of Design Tools for Quantum Computing: Arrays, Decision Diagrams, Tensor Networks, and ZX-Calculus*. In: *Proceedings of the 59th ACM/IEEE Design Automation Conference, DAC '22*, Association for Computing Machinery, New York, NY, USA, p. 1367–1370, doi:10.1145/3489517.3530627.

A Cut order correction


It was shown in section 3.3 that determining the *order* in which to cut the vertices is apparently at least as important as determining *which* vertices to cut. In the example showcased there, it seemed necessary for the optimal solution to cut vertex #8 first, such that vertices #2, #6, #11, and #15 could then be fused and cut as one. It would appear that recognising the same vertices to cut but applying a different cut ordering (specifically cutting vertex #8 *last*) would achieve the same ends at the cost of 5 cuts rather than 2 (and

hence a much less efficient 32 stabiliser terms rather than 4). However, with slight alteration (using the modified PyZX package of [27]) to denote the cuts parametrically, followed by some simple parametric analysis, a suboptimal cut ordering can effectively be corrected to its more optimal arrangement at a negligible cost to the runtime.


Firstly, one can parameterise the cutting decomposition of section 3.1 like so:

$$\begin{array}{c} \vdots \\ \textcircled{\alpha} \\ \vdots \end{array} \approx \begin{array}{c} \vdots \\ \textcircled{\color{red}\pi} \textcircled{\color{red}\pi} \\ \vdots \end{array} + e^{i\alpha} \begin{array}{c} \vdots \\ \textcircled{\color{red}\pi} \textcircled{\color{red}\pi} \\ \vdots \end{array} \equiv \sum_{a=0}^1 e^{i\alpha a} \begin{array}{c} \vdots \\ \textcircled{\color{red}a\pi} \textcircled{\color{red}a\pi} \\ \vdots \end{array}$$

Now consider again the circuit shown in section 3.3, and imagine cutting first vertices #2, #6, #11, and #15. Writing this in the parameterised form - needing only one parameterised graph (rather than 16 evaluated graphs) - leads to the following, after only some trivial spider fusion:



Cutting lastly vertex #8 results in the following:



Here we have a graph in 5 parameters (a, b, c, d, e) , containing 3 free nodes (legless spiders). As the parameters are boolean, it necessarily follows that each of these free nodes can be replaced with a scalar factor of 0 or 2, such as follows:

$$\textcircled{(a+b+e)\pi} = \begin{cases} 2 & \text{if } a \oplus b \oplus e = 0 \\ 0 & \text{if } a \oplus b \oplus e = 1 \end{cases}$$

Moreover, for any combination of parameter values that results in a scalar factor of 0 in one or more of these free nodes, one can ignore the entire corresponding graph (as the whole graph then becomes 0). In other words, one is only interested in the sets of parameter values which result in non-zero scalar factors for all free nodes. So, given the equation above, the leftmost free node must equal 2 and thus $a \oplus b \oplus e = 0$. Rearranging this for, say, a gives: $a = b \oplus e$. And now, every instance of a throughout the parameterised graph can be substituted out for $b \oplus e$, thus reducing the number of parameters from 5 to 4. Repeating this reasoning for the remaining two free nodes finds that $b = c \oplus e$ and $c = d \oplus e$, resulting in all parameters being reduced to some combination of d and e .

Consequently, while only needing to reason on one graph, the number of parameters has been reduced from the 5 attained via a suboptimal cut ordering to the most optimal 2. Indeed, this graph is now equivalent to what would have been attained if the more optimal cut ordering had been adopted (namely


cutting vertex #8 first and then fusing the remaining 4 before cutting them as one). Having effectively corrected suboptimal cut ordering, one can then expand the parameterised graph out into its (in this case 4) distinct evaluated graphs and proceed with simplification as in section 3.

Note, however, that the procedure outlined in section 3.4 prioritises higher tiers so that the cut ordering is already optimised and thus this parametric reasoning should not generally be necessary. Nevertheless, it might prove beneficial to the procedure in extreme cases, and indeed is applicable in other cutting techniques (such as the brute-force verifications of section 4.1).

B Generating random pseudo-structured circuits

The procedure presented in this paper, by design, works best on circuits that are highly structured. This raised an interesting problem when benchmarking, as testing on wholly random circuits would not properly showcase its effectiveness and, conversely, demonstrating only ideal example cases would not yield particularly informative results as such circuits could be made arbitrarily ideal, sending $\alpha \rightarrow 0$. Consequently, for the benchmarking experiments, we generated random circuits in such a way as to include some localised structural elements, so as to avoid trivial (unstructured) cases, while also not simply designing “best case” circuits on which to experiment.

Specifically, for the main benchmarking tests (figure 4), we generated circuits from random combinations of randomly placed T-CNOT-T sandwiches, Toffoli gates, CNOTs, and phase gates. We varied the number of such instances of these components in order to vary the T-counts of the circuits. Meanwhile, the circuits for the initial, small-scale verification experiments (section 4.1) were generated likewise, minus the Toffoli gates (as even one or two of these produce circuits too large to verify). Note that in randomly placing multi-qubit gates, each of its qubits may be placed randomly among the circuit’s qubits, and note also that we express Toffoli gates in ZX-diagram form as follows:



This approach ensures we create appropriately random circuits of non-trivial local structure, akin to that of section 3.3 or the composed-Toffolis circuits seen in [18].

C Estimating efficiency

Following the procedure of section 3.4 will result in a list of scalar terms whose sum gives the overall scalar value corresponding to the original Clifford+T circuit. However, for the purposes of benchmarking the effectiveness of the procedure, it is not important what that scalar value is, but rather how many scalar terms were produced (and hence the total runtime). With this in mind, one can determine an approximate upper-bound estimate of the number of terms that would be produced, in linear time, without having to compute all such terms.

This can be done by parameterising the branches of each cut into a single parametric graph (rather than two sets of distinct enumerated graphs), akin to [29] (and utilising the corresponding modified PyZX package of [27]). As such, every cut can be recorded parametrically on a single graph, without the need to append a list of distinct graphs, similar to the figures shown in appendix A. By continuing the procedure on this parametric graph, one will ultimately arrive at a final Clifford graph of n parameters (given n cuts), which represents 2^n distinct graphs whose sum is the solution. Thus, the number of terms

will be known to be 2^n . This is, however, an upper-bound estimate as the parameterisation of the cuts can prevent much inter-step simplification that could be facilitated with exact numerical phases.