

---

Cambridge Quantum

# Fast and Effective Techniques for T-Count Reduction via Spider Nest Identities

*by Niel de Beaudrap, Xiaoning Bian, Quanlong Wang*

Department of Computer Science  
University of Oxford, United Kingdom  
NIEL DE BEAUDRAP  
niel.debeaudrap@cs.ox.ac.uk

Department of Mathematics & Statistics  
Dalhousie University, Canada  
XIAONING BIAN  
bian@dal.ca

Department of Computer Science  
University of Oxford, United Kingdom  
Cambridge Quantum  
QUANLONG WANG  
quanlong.wang@cs.ox.ac.uk

Cambridge Quantum  
Terrington House  
13-15 Hills Road  
Cambridge CB2 1NL  
United Kingdom

Published by Cambridge Quantum

10 April 2020

# Fast and effective techniques for T-count reduction via spider nest identities

Niel de Beaudrap

Department of Computer Science  
University of Oxford  
Oxford, UK  
niel.debeaudrap@cs.ox.ac.uk

Xiaoning Bian

Department of Mathematics & Statistics  
Dalhousie University  
Halifax, Canada  
bian@dal.ca

Quanlong Wang

Department of Computer Science  
University of Oxford  
Oxford, UK  
Cambridge Quantum Computing Ltd.  
Cambridge, UK  
quanlong.wang@cs.ox.ac.uk

## Abstract

In fault-tolerant quantum computing systems, realising (approximately) universal quantum computation is usually described in terms of realising Clifford+T operations, which is to say a circuit of CNOT, Hadamard, and  $\pi/2$ -phase rotations, together with  $T$  operations ( $\pi/4$ -phase rotations). For many error correcting codes, fault-tolerant realisations of Clifford operations are significantly less resource-intensive than those of  $T$  gates, which motivates finding ways to realise the same transformation involving  $T$ -count (the number of  $T$  gates involved) which is as low as possible. Investigations into this problem [3–6, 13, 21] has led to observations that this problem is closely related to NP-hard tensor decomposition problems [23] and is tantamount to the difficult problem of decoding exponentially long Reed-Muller codes [6]. This problem then presents itself as one for which must be content in practise with approximate optimisation, in which one develops an array of tactics to be deployed through some pragmatic strategy. In this vein, we describe techniques to reduce the  $T$ -count, based on the effective application of “spider nest identities”: easily recognised products of parity-phase operations which are equivalent to the identity operation. We demonstrate the effectiveness of such techniques by obtaining improvements in the  $T$ -counts of a number of circuits, in run-times which are typically less than the time required to make a fresh cup of coffee.

## 1 Introduction

To achieve practical scalable quantum computation, it is important to find effective (both useful and efficient) techniques to reduce the resources required to perform computations. Error correction, and in particular realising operations in a fault-tolerant way, is expected to be a particularly significant source of resource overheads. In most quantum error-correcting codes, Clifford group operations involve less overhead than non-Clifford gates, such as the  $T$  (or  $\pi/4$  phase-rotation) gate. As the set of Clifford+T circuits is approximately universal for quantum computation [32], this motivates the  $T$ -count — or the number of  $T$  gates — as a quantity of interest in the resources required to realise a quantum computation.

On the other hand, in order to test the effectiveness of quantum technologies, it is helpful to be able to simulate the outcomes of quantum computations inasmuch as this is feasible. As circuits of Clifford operations can be efficiently simulated [1, 22], this motivates the approach of simulating quantum circuits by extending those efficient simulation techniques [11, 12], this again motivates the  $T$  count as a measure of interest in the complexity of quantum circuits.

In this article, we consider the problem of reducing the  $T$ -count required to represent a unitary circuit provided as input. Following Heyfron and Campbell [23], we consider transformations of circuits which isolate a subcircuit of diagonal operations which is the only part of the algorithm with non-trivial  $T$ -count. The approach of Heyfron and Campbell [23] is to transform Clifford+T circuits, to circuits with the following structure:

- (i) An initial stage of CNOT gates; followed by
- (ii) A stage of diagonal non-Clifford operations; followed by
- (iii) A sequence of (possibly classically controlled) Clifford operations.

This allows Ref. [23] to reduce the problem of  $T$ -count reduction to an analysis of the diagonal non-Clifford portion of this circuit, in terms of *phase polynomials*. This builds on a sequence of results which revolve around such operations [3–6, 13, 21] presented in various but similar ways, and in particular establishes a connection between  $T$ -count optimisation and difficult coding problems and tensor decomposition problems [6, 23]. Our approach is to elaborate on that of Campbell and Heyfron as follows:

- Reduce the complexity of the diagonal non-Clifford operation by more flexible (but essentially elementary) separation of the circuit into stages by allowing the first stage to contain arbitrary Clifford gates;
- Analyse the diagonal non-Clifford portion of the circuit directly in terms of “ $\pi/4$ -parity-phase operations” — essentially operators of the form  $\exp(i\frac{\pi}{8}(Z \otimes \cdots \otimes Z))$  — rather than as phase polynomials, simplifying them through the efficient application of identities of such operations.

We call these “ $\pi/4$ -parity-phase operations” as they induce a  $e^{i\pi/4}$  relative phase on standard basis states, depending on some parity computation  $f(x) = x_{k_1} \oplus x_{k_2} \oplus \cdots \oplus x_{k_m}$ . As each  $\pi/4$ -parity-phase gate can be realised in principle using a single  $T$  or  $T^\dagger$  gate (and some CNOT gates), simplifying  $\pi/4$ -parity-phase circuits is directly productive to reducing  $T$ -count.

This line of investigation, first identified in the context of  $T$ -count by Amy, Maslov, and Mosca [4], was further developed upon by Gosset *et al.* [21], Amy and Mosca [6], Kissinger and van de Wetering [26], and Zhang and Chen [34]. In previous work [7], we described a family of identities of  $\pi/4$ -parity-phase operations — “spider nest identities” — which, when used in combination with Heyfron and Campbell’s “TODD” subroutine [23], led to new records in  $T$ -count for several benchmark circuits.

In this work, we report new techniques for  $T$ -count reduction through the use of spider nest identities, and compare their effectiveness (the reduced  $T$  count and run-times) against the best previous result found in the literature. While these techniques could easily be combined with other high-performance reduction subroutines such as TODD, our results do not involve any other recently developed techniques beyond those of Ref. [7]. We obtain a number of new records for the  $T$ -count, obtained almost exclusively<sup>1</sup> in very practical run-times on a consumer-grade laptop. (For example, the second-largest circuit, on 768 qubits, was simplified in less than 3 minutes.) This opens the door to further improvements through the identification of further useful identities of  $\pi/4$ -parity-phase operations, and improved techniques for deploying these identities.

## 2 Preliminaries

We first set out some basic or existing results, using the following notation. Let  $[n] := \{1, 2, \dots, n\}$  and  $\mathbb{1}$  be the  $2 \times 2$  identity matrix. For sets  $S, T \subseteq V$  we write  $S \Delta T$  for the symmetric difference  $(S \cup T) \setminus (S \cap T)$ , and  $\mathbf{x}^{(S)} \in \{0, 1\}^V$  denote the incidence vector of  $S$ , where  $x_j^{(S)} = 1$  if and only if  $j \in S$ . We let  $\mathcal{P}^n := \{i^k P_1 \otimes \cdots \otimes P_n \mid k \in \mathbb{Z} \ \& \ P_j \in \{\mathbb{1}, X, Y, Z\}\}$  denote the  $n$ -qubit Pauli group. We define the Clifford hierarchy (on  $n$  qubits) by defining  $\mathcal{C}_1^n := \mathcal{P}_n$ , and

$$\mathcal{C}_k^n = \{U \in \mathbf{U}_n(\mathbb{C}) \mid \forall P \in \mathcal{P}^n. U P U^\dagger \in \mathcal{C}_{k-1}^n\} \quad (1)$$

<sup>1</sup>The one circuit which we did not simplify on a laptop was the largest benchmark circuit that we tested, acting on 1536 qubits and involving nearly two million  $T$  gates alone. This was instead simplified on Dalhousie University’s Mathstat Cluster [14], which took less than 15 minutes to realise a 43% reduction in  $T$ -count.

for  $k > 1$ ; we call  $\mathcal{C}_k^n$  (for arbitrary  $n$ ) the  $k^{\text{th}}$  level of the Clifford hierarchy. As an abuse of notation, we identify  $\mathcal{C}_k^n$  as a subset of  $\mathcal{C}_k^N$  for  $n < N$ ; we may then write  $S \in \mathcal{C}_2^n$  and  $T \in \mathcal{C}_3^n$  for all  $n \geq 1$ .

Let  $\mathcal{D}_k^n \subseteq \mathcal{C}_k^n$  be the subset of diagonal operations in the  $k^{\text{th}}$  level of the Clifford hierarchy. (We again identify  $\mathcal{D}_k^n$  as a subset of  $\mathcal{D}_k^N$  for  $n < N$ .) It is easy to show that  $\mathcal{D}_k^n$  forms an abelian group. In particular: consider any diagonal operation as a product of operators  $\exp(i\theta_x |x\rangle\langle x|)$  for various  $x \in \{0, 1\}^n$ , and expand each  $|x\rangle\langle x|$  as a linear combination of Pauli operators. Then one may show (see Ref. e.g. [7, Appendix A]) that  $\mathcal{D}_k^n$  is generated by the operators  $\omega \cdot \mathbb{1}^{\otimes n}$  for any global phase  $\omega$ , together with all operations of the form  $D_{S,k}$  for sets  $S = \{s_1, \dots, s_m\} \subseteq [n]$  for  $m \geq 1$ , defined by

$$D_{S,k} = \exp\left(-\frac{i\pi}{2^k}(Z_{s_1} \otimes \dots \otimes Z_{s_m})\right) = \exp\left(-\frac{i\pi}{2^k}Z_S\right) = \cos\left(\frac{\pi}{2^k}\right)\mathbb{1} - i\sin\left(\frac{\pi}{2^k}\right)Z_S, \quad (2)$$

where  $Z_S = \bigotimes_{j \in S} Z_j$ .<sup>2</sup> Note that  $X_a Z_S X_a^\dagger = (-1)^{x_a^{(S)}} Z_S$ , and that  $\text{CNOT}_{a,b} Z_S \text{CNOT}_{a,b}^\dagger = Z_{S'}$ , where here  $S' = S \Delta \{a\}$  if  $b \in S$  and  $S' = S$  otherwise. From this it follows that

$$X_b D_{S,k} X_b^\dagger = D_{S,k}^{-1} \in \mathcal{D}_k^n \quad (3a)$$

if  $b \in S$  (and  $X_b D_{S,k} X_b^\dagger = D_{S,k}$  otherwise); and

$$\text{CNOT}_{a,b} D_{S,k} \text{CNOT}_{a,b}^\dagger = D_{S',k} \in \mathcal{D}_k^n \quad (3b)$$

so that  $\mathcal{D}_k^n$  is preserved under conjugation by CNOT and  $X$  operations. Also note that  $D_{S,k}^2 = D_{S,k-1}$ , from which it follows that  $\mathcal{D}_{k-1}^n \subseteq \mathcal{D}_k^n$ .

We refer to the operators  $D_{S,k+1}$ , and their inverses, as “ $\pi/2^k$ -parity-phase” operations, as the action of  $D_{S,k+1}$  on standard basis states is given by

$$D_{S,k+1} |z\rangle = e^{i\pi/2^{k+1}} \exp\left(i[\mathbf{x}^{(S)} \cdot z] \pi/2^k\right) |z\rangle \quad (4)$$

inducing a relative phase of  $\pi/2^k$  depending on the result of a parity computation  $\mathbf{x}^{(S)} \cdot z = z_{s_1} \oplus z_{s_2} \oplus \dots \oplus z_{s_m}$ . More generally, we may refer to  $\exp(\pm \frac{1}{2} i\theta Z_S)$  as a  $\theta$ -parity-phase operation.

From Eqn. (3b), it follows that any operation  $D_{S,k}$  can be reduced to an operation  $D_{j,k} \propto \text{diag}(1, e^{2\pi i/2^k})$  acting on a single qubit  $j$ , by conjugation with an appropriate CNOT circuit. In particular, it follows that the operation  $D_{S,3}$  can be easily realised with a  $T$ -count of 1. This allows us to approach the question of reducing  $T$  count by considering decompositions of unitaries involving few  $\pi/4$ -parity-phase operations, acting on many qubits. Amy and Mosca [6] noted the relevance of the operators  $D_{S,k}$  in this context, and both Kissinger and van de Wetering [26] and Zhang and Chen [34] make direct use of them in their analysis of  $T$  count to achieve their results. (Litinski [27] similarly considers these operators in the context of compilation of quantum circuits to lattice surgery [24]).

An important role of  $D_{S,3}$  gates for  $S \subseteq [n]$  is their relationship to diagonal gates in  $\mathcal{D}_3^n$  which are controlled-unitaries in a more straightforward sense, such as CS and CCZ:

$$\text{CS} = \exp\left(\frac{i\pi}{2} |11\rangle\langle 11|\right), \quad \text{CCZ} = \exp\left(i\pi |111\rangle\langle 111|\right); \quad (5)$$

we may describe how to generate these from  $D_{k,3}$  operations by decomposing the projectors  $|11\rangle\langle 11|$  or  $|111\rangle\langle 111|$  into tensor products of  $|1\rangle\langle 1| = \frac{1}{2}(\mathbb{1} - Z)$ , and expanding to obtain a product of  $D_{S,3}$  gates. Disregarding any  $D_{\emptyset,3}$  factors, which realise global phases, we obtain

$$\text{CS}_{h,j} \propto D_{\{h\},3} D_{\{j\},3} D_{\{h,j\},3}^{-1}; \quad \text{CCZ}_{g,h,j} \propto D_{\{g\},3} D_{\{h\},3} D_{\{j\},3} D_{\{g,h\},3}^{-1} D_{\{g,j\},3}^{-1} D_{\{h,j\},3}^{-1} D_{\{g,h,j\},3}. \quad (6)$$

<sup>2</sup>We define  $D_{S,k}$  for all  $k \in \mathbb{Z}$ ; however, as one may easily show  $D_{S,0} = -\mathbb{1}^{\otimes n}$  and  $D_{S,k} = \mathbb{1}^{\otimes n}$  for all  $k < 0$  and  $S \subseteq [n]$ , these operations are of interest principally for  $k > 0$ .

More generally, we may relate  $(t-1)$ -controlled  $\pi/2^k$ -phase gates to  $\pi/2^{k-t+1}$ -phase parity gates:

$$\prod_{\substack{S \in \mathcal{P}(V) \\ S \neq \emptyset}} D_{S,k}^{(-1)^{|S|}} \propto \exp\left(\frac{i\pi}{2^{k-|V|+1}} |1\rangle\langle 1|^{\otimes V}\right), \quad (7)$$


where the right-hand operator applies a phase of  $\pi/2^{k-|T|+1}$  to those components of a state in which all of the qubits in  $T$  are in the state  $|1\rangle$ .

Circuits of parity-phase operations on  $n$  qubits which realise the identity, correspond in the notation of Amy and Mosca [6] to operators  $U_{P_{\mathbf{a}}}$  for  $\mathbf{a} \in \mathcal{C}_n \subseteq \mathbb{Z}_8^{2^n-1}$ , where

$$P_{\mathbf{a}}(\mathbf{z}) = \sum_{\substack{\mathbf{x} \in \{0,1\}^n \\ \mathbf{x} \neq \mathbf{0}}} a_x (x_1 z_1 \oplus x_2 z_2 \oplus \cdots \oplus x_n z_n) \quad (8)$$

and where  $U_{P_{\mathbf{a}}}|\mathbf{z}\rangle = \exp\left(\frac{i\pi}{4} P_{\mathbf{a}}(\mathbf{z})\right)|\mathbf{z}\rangle$ , which is identically  $|\mathbf{z}\rangle$  for all  $\mathbf{z} \in \{0,1\}^n$  when  $\mathbf{a} \in \mathcal{C}_n$ . Let  $\text{supp}(\mathbf{a}) = \{\mathbf{x} \in \{0,1\}^n : a_x \neq 0\}$ . In this notation, each element  $\mathbf{y} \in \text{supp}(\mathbf{a})$  corresponds to a single phase-parity operator acting on the qubits  $j$  for which  $y_j = 1$ ; the relative phase induced by this operator is  $a_{\mathbf{y}}\pi/4$ ; and the polynomial  $P_{\mathbf{a}}$  describes a commuting product of such operations, for which  $P_{\mathbf{a}} : \{0,1\}^n \rightarrow \mathbb{Z}_8$  is the all-zero function when  $\mathbf{a} \in \mathcal{C}_n$ .

We remark that a  $\theta$ -phase parity operation  $U$  (such as an operator  $D_{S,k}$ ) can be easily represented as tensor networks, using ZX diagrams (see Appendix A for an introduction to this notation),<sup>3</sup> with structure such as the following:



where horizontal wires represent qubits which are acted on by  $U$ , and  $S \subseteq [n]$  is the subset of those qubits which have (light, green) degree-3 nodes on them. These are “phase gadgets”, using the terminology of Kissinger and van de Wetering [26]. When the number of qubits acted on is  $m$ , we may refer to it as an “ $m$ -gadget”. (If  $\theta$  is an odd multiple of  $\pi/4$ , we may refer to it as a “ $T$ -phase  $m$ -gadget”; for  $\theta$  an integer multiple of  $\pi/2$ , we refer to it as a “Clifford-phase  $m$ -gadget”. If  $m = 1$ , we may also mildly abuse this terminology to refer to a simple green phase node as a “1-gadget”.)

**Remark.** The role played by the ZX calculus in our work is not an essential one, nor is expertise in the ZX calculus required to understand our results. However, in practice it did inform our line of investigation, by allowing us to obtain our results more quickly by identifying the objects of interest, and by making it easy to reason directly about the operators  $D_{S,k}$ . As the ZX calculus also provides a useful notation for visually representing the (non-local) unitary gates  $D_{S,k}$  in a readable way, as in Eqn. (9), we use this notation in the article below. Readers should be able to understand our results by reading ZX diagrams simply as a straightforward alternative notation for quantum circuits (see Appendix A), the transformations of which are the subject of our work.

<sup>3</sup>In this article, where they occur, ZX diagrams may be read essentially as circuit diagrams, and in particular are read from left to right as with other circuit diagrams.

### 3 Phase gadget elimination tactics & spider nest identities

Reducing the  $T$ -count while preserving the meaning of a circuit, implicitly involves applying a mathematical identity. These are often identities of diagonal unitary circuits [4, 6, 34], though not always [21, 26].) In the special case of unitary circuits consisting solely of  $\pi/4$ -parity-phase operations, such a mathematical identity may be described in terms of a commuting product of operations which are proportional to the identity operator; and for any such identity, there is the question of how to effectively apply it to realise a significant reduction of  $T$ -count, as efficiently as possible.

In this section, we describe a broad framework for the reduction of  $T$ -count by means of the application of mathematical identities of commuting  $\mathcal{D}_3^n$  operations. We also present some mathematical identities of this form — called “spider nest identities” — first presented in Ref. [7], and describe new techniques to use these identities to reduce  $T$ -count.

In the following, we use the terms “identity of  $\pi/4$ -parity-phase operations” or “identity of phase gadgets” (or simply “an identity”) to refer to a circuit  $\mathcal{J}$ , whose  $T$ -count is at least 1 but which nevertheless realises the identity operation.

#### 3.1 PHAGE tactics

We consider a particular approach to the reduction of  $\mathcal{D}_3^n$  circuits by an analysis of families of non-trivial circuits which realise the identity transformation, which may be applied more broadly than we do here (and which in principle can be used to describe some existing techniques [6, 23]). For any family  $\mathcal{F}$  of identities of  $\pi/4$ -parity-phase operations, there is an associated “phase gadget elimination tactic” (or PHAGE tactic) to reduce the  $T$ -count in a circuit  $\mathbf{C}$  of such phase gadgets:

##### PHAGE Tactic ( $\mathcal{F}$ ):

1. Determine whether there is an identity  $\mathcal{J} \in \mathcal{F}$ , such that  $\mathbf{C}$  contains at least half of the  $T$ -gadgets which occur in  $\mathcal{J}$  (or their inverses).
2. For any such identity  $\mathcal{J}$ , compute a circuit  $\mathbf{C}_{\mathcal{J}}$  as the product of  $\mathbf{C}$  and  $\mathcal{J}^{-1}$ . This may allow for simplifications (using the fact noted in Section 2 that  $D_{S,k}^2 = D_{S,k-1}$ ), where by  $T$ -gadgets accumulate to form Clifford gadgets or to cancel altogether. Determine the resulting  $T$ -count.
3. Replace  $\mathbf{C}$  with the circuit  $\mathbf{C}_{\mathcal{J}}$  with the smallest  $T$ -count, if this is less than the  $T$ -count of  $\mathbf{C}$  itself.

The behaviour of a PHAGE tactic is in a sense “greedy”, in that it selects some circuit  $\mathbf{C}_{\mathcal{J}}$  which minimises the  $T$  count after a single application, ignoring the possibility of a more complicated sequence of reductions. The main principle of a PHAGE tactic is in that it selects a way to reduce the  $T$ -count, based on the comparison of a few different applicable identities of phase gadgets from a specific family  $\mathcal{F}$ . Such a tactic can then be applied again, or followed by other such “tactics”.

In principle, the Tpar subroutine of Ref. [6], the TOOL and TODD subroutines of Ref. [23], and the results of Zhang and Chen [34] may be interpreted as algorithms to deploy PHAGE tactics, possibly more than once in sequence, and possibly with a random choice of family  $\mathcal{F}$  (and where  $\mathcal{F}$  itself may on occasion be a singleton set). This approach to  $T$ -count reduction can be distinguished from that of Kissinger and van de Wetering [26], in which phases may be reduced in unitary circuits (or more general tensor networks) which are not diagonal.


The difficulty in reducing the  $T$ -count arises from the fact that there are a very large number of identities of  $\pi/4$ -parity-phase operations, and a large number of subsets  $S \subseteq [n]$  which one may consider. As Amy and Mosca observe [6], reducing the  $T$ -count is formally equivalent to decoding a length  $2^n - 1$

punctured Reed-Muller code, in that the smallest  $T$ -count of a circuit amounts to the distance of a ciphertext to a valid codeword of such a code. However, no polynomial-time algorithms are known for the decoding problem on such codes. The difficulty is in formulating a successful *strategy* — a means of selecting an appropriately-sized family  $\mathcal{F}$  of identities to try on a particular circuit. The question is then one of having a variety of tactics which one may efficiently explore and deploy to reduce the  $T$ -count.

### 3.2 Spider nest identities

We consider PHAGE tactics arising from identities of  $\pi/4$ -parity-phase operations (*i.e.*, of  $T$ -phase gadgets) which can be composed from some specific circuits — introduced in Ref. [7], and which we call “spider nest identities” — which realise the identity operator.

In qualitative terms, a “spider nest identity” consists of any circuit of phase-parity operations which realises an operation on  $n$  qubits which is proportional to the identity, in which only “very few” operations act on “many” qubits, and the vast majority act on “very few” qubits. (In terms of the notation of Amy and Mosca [6], they would correspond to  $\mathbf{a} \in \mathbb{Z}_8^{2^n-1}$  for which only very few  $\mathbf{y} \in \text{supp}(\mathbf{a})$  have Hamming weight larger than some low threshold  $w > 0$ ; in the case of  $\mathcal{D}_3^n$  operations, we set  $w = 3$ .) We generate these circuits from a minimal family of such circuits for  $n \geq 4$ , involving a single phase 4-gadget and various phase  $k$ -gadgets with  $k \leq 3$ :



Here, the  $n$ -qubit circuit on the left-hand side of Eqn. (10) consists of a 1-gadget with phase  $(n-2)(n-3)\frac{\pi}{8}$  on each line, a 2-gadget on each pair of lines with phase  $-(n-3)\frac{\pi}{4}$ , and a 3-gadget with phase  $\frac{\pi}{4}$  on each set of three lines, and finally an  $n$ -gadget with phase angle  $-\frac{\pi}{4}$ . (For a proof of this identity, see Appendix B of Ref. [7]; in the case  $n = 4$  this corresponds to  $R_{13}$  of Ref. [3].) The name “spider nest” here refers to the qualitative feature that it involves a few “large spiders”, together with a large number of “small spiders”.

Let  $\mathcal{N}_S$  represent the circuit of phase gadgets on the left-hand side of Eqn. (10), acting on a set  $S = \{1, 2, \dots, n\}$  of cardinality  $n$ . How easily one may use this identity as part of a PHAGE tactic, to reduce  $T$ -count, is affected by the  $T$ -count of the circuit  $\mathcal{N}_S$  itself. For a fixed value of  $n$ , and a  $T$ -phase gadget on 1 to 3 qubits, there is a question of whether or not such a gadget is involved in  $\mathcal{N}_S$ , as a number of the phase gadgets involved are Clifford-phase gadgets instead. In particular:

- If  $n \equiv 1 \pmod{4}$  or  $n \equiv 3 \pmod{4}$ , all of the 2-gadgets in Eqn. (10) are Clifford-phase gadgets, which do not contribute to the  $T$ -count.
- If  $n \equiv 2 \pmod{4}$  or  $n \equiv 0 \pmod{4}$ , all of the 1-gadgets in Eqn. (10) are Clifford-phase gadgets, which again do not contribute to the  $T$ -count.


Let  $\mathbf{T}_n$  denote the  $T$ -count of  $\mathcal{N}_S$ : then

$$\mathbf{T}_n = \begin{cases} \frac{1}{6}n(n^2 + 6\delta_n - 1), & \text{for } n \text{ even;} \\ \frac{1}{6}n(n^2 - 3n + 6\delta_n + 2), & \text{for } n \text{ odd,} \end{cases} \quad (11)$$

where  $\delta_n = 1$  if  $n \equiv 0$  or  $n \equiv 1$  modulo 4, and  $\delta_n = 0$  if  $n \equiv 2$  or  $n \equiv 3$  modulo 4 (determining whether the 1-gadgets on each wire have  $T$ -count one or zero). In general, we have  $\mathbf{T}_n = \frac{1}{6}n^3 - O(n^2) \pm O(n)$ .

The scaling of  $\mathbf{T}_n$  above might suggest that these circuits have at best a limited role to play in  $T$ -count reduction: for increasing sizes of wire-sets  $S$ , a somewhat large number of operations on a given subset  $S$  of wires must be present for substitution of  $\mathcal{N}_S$  to yield a reduction in  $T$ -count. However, by composing multiple such circuits  $\mathcal{N}_S$  for different subsets  $S$ , we may obtain a ‘‘composite’’ spider nest identity which has a smaller  $T$ -count, and which is thus more likely to be usable in practise for  $T$ -count reduction.

For instance, consider the specific circuit  $\mathcal{N}_S \mathcal{N}_{S'}^{-1}$  where  $|S| \geq 5$  and  $S' = S \setminus \{r\}$  for some  $r \in S$ . As all of the operations in these circuits commute, it is possible to see that most of the phase 3-gadgets of  $\mathcal{N}_S$  — the dominant contribution to  $\mathbf{T}_n$  above — are cancelled by corresponding phase 3-gadgets of  $\mathcal{N}_{S'}^{-1}$ . (In many cases, most of the phase 1-gadgets of  $\mathcal{N}_S$  are similarly cancelled.) By collecting together the actions of the phase gadgets on each subset, we may show that  $\mathcal{N}_S \mathcal{N}_{S'}^{-1}$  simplifies to a circuit of the following form:



If  $r = S \setminus S'$  represents the top qubit in the circuit above, note in particular that the dominant contributions to the size of the circuit are the phase 2-gadgets on all size-2 subsets of  $S'$ , and the phase 3-gadgets which involve  $r$  and some size-2 subset of  $S'$ . If  $\tilde{\mathbf{T}}_n$  denotes the  $T$ -count of the circuit above, we then have

$$\tilde{\mathbf{T}}_n = \begin{cases} n^2 - n + 2 + \delta_n & \text{for } n \text{ even;} \\ n^2 - 3n + 4 + \delta_n & \text{for } n \text{ odd,} \end{cases} \quad (13)$$

where again  $\delta_n = 1$  if  $n \equiv 0$  or  $n \equiv 1$  modulo 4, and  $\delta_n = 0$  if  $n \equiv 2$  or  $n \equiv 3$  modulo 4. In any case, we have  $\tilde{\mathbf{T}}_n = n^2 - O(n)$ .

### 3.3 Simple PHAGE tactics based on spider nest identities

Combining the two ideas above, we describe the PHAGE tactics which are used to achieve the  $T$ -count reductions seen in our results.

The first tactic is the reduction of phase-parity circuits by merging together  $\pi/4$ -parity-phase operations which act on sets of qubits in common, which may be described as the PHAGE tactic associated to the circuits consisting of mutually inverse pairs of  $T$ -phase gadgets on all possible sets of qubits. To do this to greatest effect (and also as simply as possible), we first use a circuit transformation procedure along the lines of Heyfron and Campbell [23], with modifications to improve performance. (In the context of reasoning about  $T$  count in terms of  $\pi/4$ -parity-phase operations, this technique was introduced in Ref. [7].) We describe this in more detail in the following Section, which describes our  $T$ -count reduction procedure.

Our other PHAGE tactic (or tactics, as they are similar but technically numerous) are novel, and are best described in terms of the following two sets of spider-nest identities on  $N$  qubit circuits:



- The family  $\mathcal{F}_N^{(4)} = \{\mathcal{N}_S \mid S \subseteq [N] \text{ and } |S| = 4\}$ , consisting of versions of the identity of Eqn. (10) applied to all subsets of  $[N]$  of size 4
- The family

$$\mathcal{F}_N^{(5)} = \left\{ \mathcal{N}_S^{p_0} \mathcal{N}_{S_1}^{p_1} \mathcal{N}_{S_2}^{p_2} \mathcal{N}_{S_3}^{p_3} \mathcal{N}_{S_4}^{p_4} \mathcal{N}_{S_5}^{p_5} \left| \begin{array}{l} S = \{q_1, q_2, q_3, q_4, q_5\} \text{ for distinct } q_j \in [N], \\ S_j = S \setminus \{q_j\} \text{ for } 1 \leq j \leq 5, \text{ and} \\ p_0 p_1 p_2 p_3 p_4 p_5 \in \{0, 1\}^6 \setminus \{000000\} \end{array} \right. \right\}, \quad (14)$$

consisting of the 63 distinct identities for each set  $S \subseteq [N]$  with  $|S| = 5$ , consisting of  $\mathcal{N}_{S_j}$  applied to some or all subsets  $S_j \subseteq S$  of size 4, and possibly also a copy of  $\mathcal{N}_S$  on all the qubits of  $S$ , fusing together those phase-parity operations which act on common subsets  $S' \subseteq S$ .

These are the sets of all possible spider-nest identities on 4 or 5 qubits.<sup>4</sup>

For increasing values of  $N$ , the cardinalities of these families grow as  $\frac{1}{24}n^4 + O(n^3)$  and  $\frac{1}{120}n^5 + O(n^4)$  respectively — polynomial in size, but impractical to exhaustively iterate through for values of  $N$  which occur in common benchmark tests. This raises the question of how best to use them to realise  $T$ -count reductions. Our approach is to construct a list of 64 identities on four or five qubits, consisting of the elements of the sets  $\mathcal{F}_4^{(4)} \cup \mathcal{F}_5^{(5)}$ , and performing the following for each element  $\mathcal{J}$  of this list:

1. Let  $s$  be the number of qubits on which  $\mathcal{J}$  acts.
2. Repeat the following  $R$  times, for some fixed  $R > 0$ :
  - a. Select a subset  $S \subseteq [N]$  of size  $s$  uniformly at random.
  - b. Select (from  $\mathcal{F}_N^{(4)}$  if  $s = 4$ , or  $\mathcal{F}_N^{(5)}$  if  $s = 5$ ) the identity  $\mathcal{K}$  acting on  $S$ , which is equivalent to  $\mathcal{J}$  up to relabelling of the qubits.
  - c. Apply the tactic **PHAGE**( $\{\mathcal{K}\}$ ) associated with the singleton set  $\{\mathcal{K}\}$ .

This technique implicitly provides opportunities for identities to be applied in proportion to the number of isomorphic images of it exist in  $\mathcal{F}_4^{(4)} \cup \mathcal{F}_5^{(5)}$ . (For instance, isomorphic copies of the simplest identity  $\mathcal{N}_{[4]}$  occurs six times in this set, and the identity of Eqn. (12) occurs five times.) As the probability that any one such identity will be useful when applied to a particular set  $S \subseteq [N]$  of size 4 or 5 is small, it is important to choose a significantly large value of  $R$ : for our results, we took  $R = 20\,000$ .

We note that this particular strategy for  $T$ -count reduction is not particularly strongly suggested by the framework of PHAGE tactics induced by spider nest identities. Both the concept of a PHAGE tactic, and the range of possibilities for assembling spider nest identities, are broad enough that there is potential for much more sophisticated strategies to deploy them. Despite this, as we show in Section 5, in many cases we obtain the best known  $T$ -count for a number of circuits. Our result may therefore be considered a further proof of principle of the usefulness of spider nest identities, beyond the results of Ref. [7].

## 4 Reduction of $T$ -count through simplification of parity-phase circuits

In this section, we describe how we applied the concept of  $T$ -count reduction via PHAGE tactics as part of a complete procedure to transform unitary circuits provided as input.

<sup>4</sup>The set  $\mathcal{F}_N^{(5)}$  in particular is motivated by the reduction in  $T$ -count of the spider-nest identity shown in Eqn. (12), which is represented in five different ways in  $\mathcal{F}_N^{(5)}$ : once for each subset  $S_j$  of size 4.

**Remark.** Our results do not make heavy (explicit) use of the re-write rules of the ZX calculus: a reader who is content with circuits which involve intermediate measurements, and who is comfortable with reading a parity-phase gadget such as that of Eqn. (9) as a unitary operator, may interpret every diagram below as a circuit diagram. (See Appendix A for a guide to reading ZX diagrams.)

We take as input a unitary circuit over the gate-set  $\{X, \text{CNOT}, \text{CCNOT}, Z, \text{CZ}, \text{CCZ}, H, S, T, \text{SWAP}\}$ . For the sake of simplicity, we suppose that any multiply-controlled NOT gates with more than two controls are decomposed into CCNOT gates, for instance by computation and uncomputation on auxiliary qubits initialised to  $|0\rangle$ , or some more advanced technique.<sup>5</sup>

Our procedure follows and extends the approach of Heyfron and Campbell [23], of performing a transformation on circuits  $\mathbf{C} \rightarrow \mathbf{C}_F \circ \mathbf{C}_\phi \circ \mathbf{C}_I$ , where  $\mathbf{C}_F$  and  $\mathbf{C}_I$  consist entirely of Clifford gates, stabiliser state preparations, and stabiliser state measurements, and where  $\mathbf{C}_\phi$  can be realised using only CNOT and  $T$  gates. We express the circuit  $\mathbf{C}_\phi$  entirely in terms of phase gadgets, and so we describe as a ‘‘homogeneous’’ circuit. The objective of isolating such a circuit is that it provides us with the best opportunities to apply PHAGE tactics to reduce the  $T$ -count.

### 4.1 Circuit translation techniques

Our procedure, which we describe more explicitly in the next section, makes use of the following techniques.

**$H$  gate gadgetisation.** One of the techniques involved in isolating a  $\mathcal{D}_3^N$  circuit is to replace Hadamard gates with a measurement-based gadget:

The diagram shows the equivalence between a Hadamard gate and a measurement-based gadget. On the left, a Hadamard gate symbol is shown. This is followed by an equivalence symbol and a circuit diagram where a qubit in state  $|+\rangle$  is entangled with another qubit, which is then measured in the  $X$  basis. This is further equivalent to a ZX diagram with two qubits, each having a  $-\pi/2$  phase, and a phase gadget with a  $\pi$  phase and a heralded outcome  $s$ . The final ZX diagram shows the decomposition of the CZ operation into single-qubit  $D$  operations and a phase gadget.

In the circuit second from the left, the two qubits are subject to a SWAP operation, followed by a  $\text{CZ} = \exp(i\pi|11\rangle\langle 11|)$  operation. The bottom qubit is measured finally with an  $X$  observable measurement (*i.e.*, in the  $|\pm\rangle$  basis), and the top operation is acted on finally by an  $X$  operation only if the outcome is  $|-\rangle$ . The two diagrams on the right are ZX diagrams with additional annotations in the style of Ref. [18] (see also Appendix A). In particular, measurement is represented as a projection with a random outcome  $s$  which is heralded and may be used to control phase operations elsewhere. The leftmost ZX diagram describes the decomposition of the controlled- $Z$  operation, using  $\text{CZ}_{h,j} \propto D_{\{h,j\},2} D_{\{h\},2}^{-1} D_{\{j\},2}^{-1}$ . The final ZX diagram propagates the single-qubit  $D_{\{*\},2}^{-1}$  operations towards the preparation and measurement of the second qubit, so that the second qubit is prepared in the  $|-\rangle \propto |0\rangle - i|1\rangle$  state.

**Extracting  $H$  gates from the circuit.** An obvious drawback of gadgetising  $H$  gates in this way is that it requires the use of auxiliary qubits. More directly important to our results is that, as the number of wires in a circuit increases, the more difficult it may be to successfully find opportunities to reduce the  $T$  count. Therefore, we attempt to transform the circuit in such a way that reduces the number of  $H$  gates from the part of the circuit with non-trivial  $T$ -count. This motivates us to define a subroutine `moveH` (which we describe at a high level in Appendix B), which transforms a circuit  $\mathbf{C}$  over our gate-set, into a pair of circuits  $(\mathbf{C}_F, \mathbf{C}')$ , obtained by attempting to commute as many Hadamard gates of  $\mathbf{C}$  to the end of the circuit as possible.

<sup>5</sup>In our benchmarks, we consider the simple computation-uncomputation approach; other techniques (see *e.g.* Refs. [20, 25, 29]) are advisable in serious production work for optimising  $T$ -count.

- We define  $(\mathbf{C}_F, \mathbf{C}') = \text{moveH}(\mathbf{C})$  in such a way that  $\mathbf{C}_F \circ \mathbf{C}' \cong \mathbf{C}$  realises the same unitary,  $\mathbf{C}_F$  contains only Clifford gates,  $\mathbf{C}'$  contains no CCNOT gates, and where the total number of Hadamard gates in  $(\mathbf{C}_F \circ \mathbf{C}')$  is at most the number of Hadamard gates in  $\mathbf{C}$ .
- We may use  $\text{moveH}$  twice, to attempt to extract Hadamard gates either from the end of the circuit  $\mathbf{C}$ , and also the beginning of the circuit  $\mathbf{C}$ . If we compute

$$(\mathbf{C}_F, \mathbf{C}') = \text{moveH}(\mathbf{C}); \quad (\tilde{\mathbf{C}}_I, \tilde{\mathbf{C}}_M) = \text{moveH}((\mathbf{C}')^{-1}); \quad (\mathbf{C}_I, \mathbf{C}_M) = (\tilde{\mathbf{C}}_I^{-1}, \tilde{\mathbf{C}}_M^{-1}), \quad (16)$$

then  $(\mathbf{C}_F \circ \mathbf{C}_M \circ \mathbf{C}_I) \cong \mathbf{C}$ , the number of Hadamard gates in  $(\mathbf{C}_F \circ \mathbf{C}_M \circ \mathbf{C}_I)$  is at most the number of Hadamard gates in  $\mathbf{C}$ , and  $\mathbf{C}_I$  and  $\mathbf{C}_F$  only contain Clifford gates.

We call  $\mathbf{C}_I$  and  $\mathbf{C}_F$  the initial and final Clifford stages of the circuit, respectively, and  $\mathbf{C}_M$  the main body of the circuit. We use this tripartite decomposition to allow us to condense the part of the circuit with non-trivial  $T$ -count in the main body, and to remove Clifford gates ( $H$  gates in particular) to the initial and final Clifford phases to the extent that this is possible.


**Phase-gadgetisation.** Through appropriate substitution of  $H$  gates by gadgets as in Eqn. (15), and substitution of CCZ with  $\pi/4$ -parity-phase operations as in Eqn. (6), we may transform the main body of the circuit so that it only contains SWAP gate,  $X$  gates, CNOT gates, CZ gates, and various phase gadgets (including powers of the  $T$  gate). We wish to transform this into a circuit consisting only of phase gadgets, by commuting everything apart from phase gadgets either to the beginning of the main body (and then removing it to the initial Clifford phase) or to the end of the main body (and then removing it to the final Clifford phase). In particular, we commute all SWAP, measurement, and  $X$  operations to the end of the circuit; we commute all preparation operations to the beginning of the circuit; and we commute each CNOT operation either to the beginning or the end according to a simple heuristic (described in Appendix B). This may transform various  $D_{S,t}$  gates by Eqns. (3), changing the set  $S$  involved and/or negating the phase, according to the following commutation relations:

$$(17)$$

$$(18)$$

$$(19)$$

**Phase gadget fusion.** A final simplifying technique is to simply multiply together any phase gadgets acting on the same set  $S$  of qubits:



In some cases, this will reduce the  $T$  count by turning two gadgets with phases  $\alpha = \frac{1}{4}k_1\pi$  and  $\beta = \frac{1}{4}k_2\pi$  (for  $k_1$  and  $k_2$  odd) into a single gadget with phase  $\alpha + \beta = \frac{1}{4}(k_1 + k_2)\pi$ , where  $k_1 + k_2$  is even.

## 4.2 Circuit translation procedure

Given a unitary circuit  $\mathbf{C}$  over the gate-set  $\{X, \text{CNOT}, \text{CCNOT}, Z, \text{CZ}, \text{CCZ}, H, S, T, \text{SWAP}\}$ , we transform  $\mathbf{C}$  as follows:

1. We first replace CCNOT operations in  $\mathbf{C}$  with  $(\mathbb{1} \otimes \mathbb{1} \otimes H) \text{CCZ} (\mathbb{1} \otimes \mathbb{1} \otimes H)$ , yielding a circuit  $\mathbf{C}'$ .
2. Transform  $\mathbf{C}' \rightarrow \mathbf{C}'_F \circ \mathbf{C}'_M \circ \mathbf{C}'_I$ , with an initial Clifford stage  $\mathbf{C}'_I$ , a final Clifford stage  $\mathbf{C}'_F$ , and a main body  $\mathbf{C}'_M$ , using the procedure `moveH` to reduce the number of Hadamard gates in  $\mathbf{C}'_M$  as much as possible.
3. Substitute the  $H$  gates in  $\mathbf{C}'_M$  with Hadamard gadgets as in Eqn. (15), using a fresh bit label for each measurement outcome; and decompose CCZ operations in  $\mathbf{C}$  using the formula of Eqn. (6), and represent  $T$  gates (on some qubit  $j$ ) by  $D_{\{j\},3}$ . Call the resulting circuit  $\mathbf{C}_M$ .
4. We gadgetize  $\mathbf{C}_M$  by commuting all gates which are not single-qubit phase gates or phase gadgets to the beginning or the end, removing these to the initial or final Clifford stages. This will generally add some number of measurements, and classically-conditioned Clifford operations, to the final Clifford stage, and some qubit preparations to the initial Clifford stage. This realises a transformation of circuits  $\mathbf{C}'_F \circ \mathbf{C}_M \circ \mathbf{C}'_I \rightarrow \mathbf{C}_F \circ \mathbf{C}'_\phi \circ \mathbf{C}_I$ .
5. As  $\mathbf{C}'_\phi$  is now a homogeneous circuit of phase gadgets, we may commute them past one another to fuse gadgets on common subsets, yielding a circuit  $\mathbf{C}_\phi$ .
6. Apply the randomised procedure for applying PHAGE tactics based on spider nest identities described in Section 3.3.

Steps 1–5 realise a transformation  $\mathbf{C} \rightarrow \mathbf{C}_F \circ \mathbf{C}_\phi \circ \mathbf{C}_I$ . If the original circuit  $\mathbf{C}$  acted on  $n$  qubits and had  $m$  Hadamard gates, then the number of Hadamard gates in  $\mathbf{C}'_M$  which are replaced in Step 3 is some  $\delta n \leq m$ . Then the circuits  $\mathbf{C}_I$ ,  $\mathbf{C}_\phi$ , and  $\mathbf{C}_F$  all act on  $N = n + \delta n$  qubits, and  $\mathbf{C}_F$  has internal structure

$$\mathbf{C}_F = \tilde{\mathbf{C}}_F \mathbf{D}_{\delta n} \cdots \mathbf{D}_2 \mathbf{D}_1, \tag{21}$$

where  $\tilde{\mathbf{C}}_F$  is some general Clifford circuit, and the circuits  $\mathbf{D}_j$  (for  $1 \leq j \leq \delta n$ ) consist of the  $j^{\text{th}}$  measurement in the  $|\pm\rangle$  basis with outcome  $s_j$  (denoted in ZX notation by a light green “ $\pi, \{s_j\}$ ” node), followed by  $\mathcal{D}_k^N$  operations conditioned on the outcome  $s_j$ .

In some instances, we find a significant reduction in the  $T$ -count simply from the fusion of phase gadgets in Step 5 of this transformation. These improvements are similar to those seen in Refs. [26, 34]. However, the purpose of this circuit transformation (as Ref. [23]) is to isolate a circuit  $\mathbf{C}_\phi$  consisting entirely of  $\mathcal{D}_3^N$  operations for some  $N$ , on which we can apply the PHAGE tactic of Step 6.

Note that  $\delta n$ , the number of additional “auxiliary” qubits involved in the circuit, is bounded above by how many Hadamard gates are either involved in  $\mathbf{C}$  or are introduced from the decomposition of CCNOT gates. More precisely, it depends on how many of these gates can be commuted from the “main body” of  $\mathbf{C}$  to the initial or final Clifford stages. For a circuit consisting of  $M$  gates, a bound for  $N = n + \delta n$  which is substantially better than  $N \leq n + M$  will be difficult to obtain, without some knowledge of the structure of  $\mathbf{C}$ . In several cases, we find that many or all of these Hadamard gates can be eliminated from the main body of the circuit: so,  $N \leq n + M$  is likely a loose upper bound in a large number of practical examples.

The largest contributions to the asymptotic run-time of the procedure above are the complexity of `moveH` in Step 2; the cumulated complexity of computing the heuristic for moving Clifford gates out of the main body of the circuit in Step 4; and the complexity of performing a PHAGE tactics in Steps 5 and 6. For  $M$  the number of gates in the input circuit, the procedure `moveH` and the procedure to commute CNOT gates out the main body take time  $O(M^2)$ , essentially due to repeatedly commuting individual gates past  $O(M)$  other gates (or computing the potential cost of doing so, in the case of the heuristic used for determining the direction to move CNOT gates). We use a hash table to store homogeneous circuits, allowing essentially  $O(1)$  time to look up the phase associated with a phase gadget acting on a particular subset (which we set to 0 when no such phase gadget is present). In Step 5, fusing together pairs of phase gadgets can be made a part of initialising this hash table, and so takes time  $O(M)$ . In Step 6, applying a PHAGE tactic associated with some given identity  $\mathcal{H}$  (which acts on at most 5 qubits) takes time  $O(1)$ ; performing this for each of the 64 identities in  $\mathcal{F}_4^{(4)} \cup \mathcal{F}_5^{(5)}$  on  $R$  uniformly random subsets takes time  $O(R) = O(1)$ , for  $R$  independent of  $M$ . Thus our procedure runs in time  $O(M^2)$ .

## 5 Results

We realised our techniques in Haskell code [10]. All but two of the circuits were obtained from Ref. [30]: the circuits “GF( $2^{256}$ ) Mult” and “GF( $2^{512}$ ) Mult” were obtained instead from Ref. [28]. With one exception, we ran our code for these benchmarks on a 2.5 GHz Intel Core i7 processor and 8 GB of 1867 MHz LPDDR3 memory, running Ubuntu Linux 18.04.4. The largest single benchmark circuit, “GF( $2^{512}$ ) Mult”, was instead reduced on Dalhousie University’s Mathstat Cluster [14]. The results are presented as Table 1, including comparisons to the best known reductions for which recorded times are available.<sup>6</sup>

Our results do not include an account of the cost of the Clifford group operations. These are also of interest in principle, though these will likely be significantly less expensive than  $T$  gates in the error-corrected setting in which the  $T$ -count becomes a meaningful quantity to reduce. We also do not describe the  $T$ -depth of our circuits, which may also be independently optimised from the  $T$ -count itself [4].

The circuits which were obtained using our techniques may be found on GitHub [10]. As our main aim was to consider reductions in  $T$ -count, our algorithm ignores the possibility that the measurement outcomes on the auxiliary qubits could be anything but  $|+\rangle$ : in the event of a  $|-\rangle$  outcome, additional Clifford group operations would be required, which however would not affect the  $T$ -count. We verified our circuits using `feynver` [2], which was recently extended to accomodate circuits involving post-selection of  $|+\rangle$  states on qubits which are maximally entangled with a set of other qubits.

---

<sup>6</sup>We do not present the best known  $T$  counts which do *not* have recorded times. We do note that for two of our results (for the circuits `Mod Red21` and `RC Adder6`) which improve on the known timed results, there are recorded untimed results which are still better: these may be found in Ref. [7].

Circuit	# qubits			$T$ count & optimisation						
	$n$ input	$\delta n$ [23]	$\delta n$ (ours)	init. # $T$	final # $T$ (prev. opt.)	Ref.	time (s)	final # $T$ (our results)	time (s)	Verified? (feynver)
Adder <sub>8</sub>	24	71	41	399	212 <sup>(a)</sup>	[23]	227.81	<b>176*</b>	24.62	✓
Barenco Tof <sub>3</sub>	5	3	3	28	14 <sup>(b)</sup>	[23]	0.01*	<b>13*</b>	0.07607	✓
Barenco Tof <sub>4</sub>	7	7	7	56	24	[23]	0.45	25	1.884	✓
Barenco Tof <sub>5</sub>	9	11	11	84	34	[23]	1.94	37	13.76	✓
Barenco Tof <sub>10</sub>	19	31	31	224	84	[23]	460.33	97	24.49	✓
CSLA MUX <sub>3</sub>	15	17	6	70	40 <sup>(b)</sup>	[23]	3.73	44	18.01	✓
CSUM MUX <sub>9</sub>	30	12	12	196	74 <sup>(a)</sup>	[23]	36.57	84	23.98	✓
GF(2 <sup>4</sup> ) Mult	12	7	0	112	56 <sup>(b)</sup>	[23]	0.55	<b>53*</b>	0.8180	✓
GF(2 <sup>5</sup> ) Mult	15	9	0	175	90 <sup>(b)</sup>	[23]	6.96	<b>88*</b>	4.279	✓
GF(2 <sup>6</sup> ) Mult	18	11	0	252	132 <sup>(b)</sup>	[23]	121.16	<b>128*</b>	7.894	✓
GF(2 <sup>7</sup> ) Mult	21	13	0	343	185 <sup>(a)</sup>	[23]	153.75	<b>167*</b>	27.21	✓
GF(2 <sup>8</sup> ) Mult	24	15	0	448	216 <sup>(a)</sup>	[23]	517.63	229	95.63	✓
GF(2 <sup>9</sup> ) Mult	27	17	0	567	295	[23]	3213.53	306	24.79	✓
GF(2 <sup>10</sup> ) Mult	30	19	0	700	351	[23]	23969.1	357	24.65	✓
GF(2 <sup>16</sup> ) Mult	48	31	0	1 792	922	[23]	76312.5	972	25.65	✓ <sup>(d)</sup>
GF(2 <sup>32</sup> ) Mult	96	–	0	7 168	4 128	[31]	1.834	<b>3 936*</b>	26.07	✓ <sup>(d)</sup>
GF(2 <sup>64</sup> ) Mult	192	–	0	28 672	16 448	[31]	58.341	<b>15 865*</b>	29.73	–
GF(2 <sup>128</sup> ) Mult	384	–	0	114 688	65 664	[31]	1744.746	<b>64 461*</b>	48.78	–
GF(2 <sup>131</sup> ) Mult	393	–	0	120 127	69 037	[31]	1953.353	<b>67 772*</b>	53.39	–
GF(2 <sup>163</sup> ) Mult	489	–	0	185 983	106 765	[31]	4955.927	<b>105 182*</b>	66.27	–
GF(2 <sup>256</sup> ) Mult	768	–	0	458 752	–	–	–	<b>260 539*</b>	137.1	–
GF(2 <sup>512</sup> ) Mult	1536	–	0	1 835 008	–	–	–	<b>1 046 964*</b>	850.7 <sup>(d)</sup>	–
Mod5 <sub>4</sub>	5	6	0	28	16 <sup>(b)</sup>	[31]	0.001*	<b>7*</b>	0.00899	✓
Mod Adder <sub>1024</sub>	28	≤ 270 <sup>(c)</sup>	304	1 995	978	[23]	665.5	1 010	27.56	✓ <sup>(d)</sup>
Mod Mult <sub>55</sub>	9	10	3	49	28 <sup>(a)</sup>	[23]	0.02	<b>19*</b>	0.5775	✓
Mod Red <sub>21</sub>	11	17	17	119	69 <sup>(b)</sup>	[23]	0.59	<b>65</b>	27.68	✓
QCLA Adder <sub>10</sub>	36	28	25	238	157	[23]	366.1	<b>147*</b>	24.96	✓
QCLA Com <sub>7</sub>	24	19	18	203	81	[23]	170.77	84	24.21	✓
QCLA Mod <sub>7</sub>	26	58	58	413	221 <sup>(a)</sup>	[23]	289.77	233	24.26	✓ <sup>(d)</sup>
RC Adder <sub>6</sub>	14	21	10	77	45 <sup>(b)</sup>	[23]	0.97	<b>38</b>	30.70	✓
NC Toff <sub>3</sub>	5	2	2	21	13	[23]	0.01*	<b>13*</b>	0.04005	✓
NC Toff <sub>4</sub>	7	4	4	35	19	[23]	0.06	<b>19*</b>	0.5322	✓
NC Toff <sub>5</sub>	9	11	6	49	25	[23]	0.4	26	2.910	✓
NC Toff <sub>10</sub>	19	16	16	119	55	[23]	44.78	60	28.01	✓
VBE Adder <sub>3</sub>	10	4	4	70	20	[23]	0.15	<b>20*</b>	1.896	✓

Table 1: Comparison of our techniques to previously reported results. • In each case, “prev. opt.” represents the best  $T$ -count with a time record (an asterisk indicates that the recorded time is an upper bound). For some circuits, better reductions without times have been reported: those indicated by <sup>(a)</sup> have a better reduction reported in Ref. [26], and those indicated by <sup>(b)</sup> have a better reduction reported in Ref. [7]. Where it was feasible to verify the correctness of our reduction with `feynver`, this is indicated; in all other cases the verification was too computationally expensive to carry out. • In each case, we also compare the number  $\delta n$  of additional “auxiliary” qubits required by our decomposition, to that of Ref. [23] (where results are available); in the case of <sup>(c)</sup>, we may only infer an upper bound on the number of auxiliary qubits used by Ref. [23]. • In our results, those  $T$ -counts which are indicated in bold are those which reproduce or surpass the  $T$ -count of the best previously known timed result. Those with an asterisk also match or surpass the best previously known untimed result. • All results of Ref. [23] were obtained on the University of Sheffield’s Iceberg HPC cluster. All results of Ref. [31] were obtained on a machine with a 2.9GHz Intel Core i5 processor and 8GB of 1867MHz DDR3 memory, running OS X El Capitan. All of our results were obtained on a machine with a 2.5GHz Intel Core i7 processor and 8GB of 1867MHz LPDDR3 memory, running Ubuntu Linux 18.04.4 — except those indicated by <sup>(d)</sup>, which were obtained on Dalhousie University’s Mathstat Cluster [14].

## 6 Discussion

Our results show that our techniques, simple as they are, are competitive with the best known techniques for reducing  $T$  count. We expect that better results should be achievable by a more refined approach to using these concepts, within the more general framework which we have described of deploying PHAGE tactics. It is not clear which further ideas may prove useful, however. For instance, in experiments for how we might choose subsets to apply PHAGE tactics to, we found that it was not helpful to bias the sets of qubits towards those qubits which were acted on by many  $T$ -phase gadgets. More work will be required to find effective ways to bias or to narrow down the ways in which spider nest identities are used to simplify homogeneous circuits.

It is remarkable that the run-times for our results in Table 1 are not more varied. Over half of our results were obtained in an amount of time between 1 and 100 seconds, for circuits over which other leading techniques [23, 31] had times which ranged over more than six orders of magnitude. Indeed, in our tests on larger circuits (and in line with the asymptotic analysis of Section 4.2), we found that the most computationally expensive part of our procedure was the relatively mundane `moveH` and `CNOT-commutation` subroutines. Refining these elementary steps may provide yet further gains. Expanding the complexity of the subroutines to apply PHAGE tactics may also yield further gains without substantial increases in run-time.

We note an optimisation problem of interest is motivated by gadgetizing Hadamard gates as in Step 3. Simply put: given an  $n$ -qubit circuit with  $M$  gates over the gate-set  $\{X, Z, S, \text{CNOT}, \text{CZ}, T, \text{CCZ}\}$ , to obtain an equivalent (unitary) circuit with the minimum number of  $H$  gates in between the first and the last non-Clifford gate.<sup>7</sup> Should this problem be solvable in  $O(M^2 \text{poly log}(M))$  time, this may further contribute to the effectiveness of our approach to  $T$ -count reduction.

Finally, we remark that while the benchmarks which we have tested have become a commonplace standard for the evaluation of such techniques, they consist entirely of circuits to realise permutation operations which would not in themselves be difficult to realise classically. (Some of these, such as the “GF( $2^n$ ) Mult” series, may be motivated on the grounds of cryptography; albeit this motivation may become less important if standard cryptographic practise incorporates post-quantum cryptography.) A larger range of circuits, including ones are motivated by the more likely practical applications of fault-tolerant quantum computation, should be of general interest for future benchmark tests.

### Acknowledgements.

N. de Beaudrap was supported in part by a Fellowship funded by a gift from Tencent Holdings (tencent.com), and by the EPSRC National Hub in Networked Quantum Information Technologies (NQIT.org). X. Bian is supported by NSERC and by AFOSR under Award No. FA9550-15-1-0331. Q. Wang is supported by Cambridge Quantum Computing Ltd. and by the AFOSR grant FA2386-18-1-4028. Our results were made possible in part by the use of the Dalhousie University Mathstat Cluster [14].

We thank Earl Campbell, Luke Heyfron, Alexander Cowtan, Aleks Kissinger, and John van de Wetering for helpful discussions. We extend a very special thanks to Matthew Amy, who wrote a small extension of `feynver` [2] to allow verification of procedures which post-select the  $|+\rangle$  state, for the express purpose of helping us to independently verify the correctness of reductions such as appear in this work and in Ref. [7]. X. Bian would like to thank his Ph.D. supervisor Peter Selinger for his support.

---

<sup>7</sup>It seems plausible that this problem would remain equally difficult without CCZ gates.

## References

- [1] S. Aaronson & D. Gottesman (2004): *Improved simulation of stabilizer circuits*. *Phys. Rev. A* 70, p. 052328.
- [2] Matthew Amy (2018): *Towards Large-scale Functional Verification of Universal Quantum Circuits*. In: *Proceedings of QPL 2018*, pp. 1–21, doi:10.4204/EPTCS.287.1. [arXiv:1901.09476]; see also [<https://github.com/meamy/feynman>].
- [3] Matthew Amy, Jianxin Chen & Neil J. Ross (2018): *A Finite Presentation of CNOT-Dihedral Operators*. *Electronic Proceedings in Theoretical Computer Science* 266, pp. 84–97, doi:10.1007/978-3-642-12821-9\_4. [arXiv:1701.00140].
- [4] Matthew Amy, Dmitri Maslov & Michele Mosca (2014): *Polynomial-Time T-Depth Optimization of Clifford+T Circuits Via Matroid Partitioning*. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 33(10), pp. 1476–1489, doi:10.1109/TCAD.2014.2341953. [arXiv:1303.2042].
- [5] Matthew Amy, Dmitri Maslov, Michele Mosca & Martin Roetteler (2013): *A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits*. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 32(6), pp. 818–830, doi:10.1109/TCAD.2013.2244643. [arXiv:1206.0758].
- [6] Matthew Amy & Michele Mosca (2019): *T-count optimization and Reed-Muller codes*. *IEEE Transactions on Information Theory* 65(8), pp. 4771–4784, doi:10.1109/TIT.2019.2906374. [arXiv:1601.07363].
- [7] Niel de Beaudrap, Xiaoning Bian & Quanlong Wang (2019): *Techniques to reduce  $\pi/4$ -parity phase circuits, motivated by the ZX calculus*. In: *to appear in Proceedings of QPL 2019*. [arXiv:1911.09039].
- [8] Niel de Beaudrap, Ross Duncan, Dominic Horsman & Simon Perdrix (2019): *Pauli Fusion: a computational model to realise quantum transformations from ZX terms*. In: *Proceedings of QPL 2019*, p. to appear. [arXiv:1904.12817].
- [9] Niel de Beaudrap & Dominic Horsman (2020): *The ZX calculus is a language for surface code lattice surgery*. *Quantum* 4, p. 218, doi:10.22331/q-2020-01-09-218. [arXiv:1704.08670].
- [10] Xiaoning Bian: “*STOMP-code*” *Github*. Available at <https://github.com/onestruggler/stomp-code/tree/8df4f46228c2f413e0cf5f8b6d25c20b6460fc0e>.
- [11] 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. Available at <https://doi.org/10.22331/q-2019-09-02-181>. [arXiv:1808.00128].
- [12] Sergey Bravyi & David Gosset (2016): *Improved classical simulation of quantum circuits dominated by Clifford gates*. *Physical Review Letters* 116, p. 250501, doi:10.1103/PhysRevLett.116.250501. [arXiv:1601.07601].
- [13] Earl T. Campbell & Mark Howard (2017): *A unified framework for magic state distillation and multi-qubit gate-synthesis with reduced resource cost*. *Physical Review A* 95, p. 022316, doi:10.1103/PhysRevA.86.022316. [arXiv:1606.01904].
- [14] Dalhousie University Mathstat Cluster: Available at <https://www.mathstat.dal.ca/cluster/doku.php>.
- [15] 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. [arXiv:0906.4725].
- [16] Bob Coecke & Aleks Kissinger (2017): *Picturing Quantum Processes: A First Course in Quantum Theory and Diagrammatic Reasoning*. Cambridge University Press, doi:10.1017/9781316219317.
- [17] Ross Duncan, Aleks Kissinger, Simon Perdrix & John van de Wetering (2019): *Graph-theoretic Simplification of Quantum Circuits with the ZX-calculus*. [arXiv:1902.03178].



- [18] Ross Duncan & Simon Perdrix (2010): *Rewriting Measurement-Based Quantum Computations with Generalised Flow*. In Samson Abramsky, Cyril Gavioille, Claude Kirchner, Friedhelm Meyer auf der Heide & Paul G. Spirakis, editors: *Automata, Languages and Programming*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 285–296, doi:10.1007/s10472-009-9141-x.
- [19] Simon Perdrix Emmanuel Jeandel & Renaud Vilmart (2019): *Completeness of the ZX-Calculus*. [arXiv:1903.06035].
- [20] Craig Gidney (2018): *Halving the cost of quantum addition*. *Quantum* 2, p. 74, doi:10.1007/s11128-011-0297-z. [arXiv:1709.06648].
- [21] David Gosset, Vadym Kliuchnikov, Michele Mosca & Vincent Russo (2014): *An Algorithm for the T-count*. *Quantum Info. Comput.* 14(15–16), pp. 1261–1276. [arXiv:1308.4134].
- [22] D. Gottesman (1997): *Stabilizer codes and quantum error correction*. Ph.D thesis.
- [23] Luke E. Heyfron & Earl T. Campbell (2018): *An efficient quantum compiler that reduces T count*. *Quantum Science and Technology* 4(1), p. 015004, doi:10.1038/srep01939. [arXiv:1712.01557].
- [24] C. Horsman, A. G Fowler, S. Devitt & R. Van Meter (2012): *Surface code quantum computing by lattice surgery*. *New Journal of Physics* 14(12), p. 123011. [arXiv:1111.4022].
- [25] Cody Jones (2013): *Low-overhead constructions for the fault-tolerant Toffoli gate*. *Phys. Rev. A* 87, p. 022328, doi:10.1103/PhysRevA.87.022328. [arXiv:1212.5069].
- [26] Aleks Kissinger & John van de Wetering (2019): *Reducing T-count with the ZX-calculus*. [arXiv:1903.10477].
- [27] Daniel Litinski (2019): *A Game of Surface Codes: Large-Scale Quantum Computing with Lattice Surgery*. *Quantum* 3, p. 128, doi:10.1103/PhysRevB.96.205413. [arXiv:1808.02892].
- [28] Dmitri Maslov: *Reversible Logic Synthesis Benchmarks page*. Available at <http://webhome.cs.uvic.ca/~dmaslov>. Accessed February 2020.
- [29] Giulia Meuli, Mathias Soeken, Earl Campbell, Martin Roetteler & Giovanni De Micheli (2019): *The Role of Multiplicative Complexity in Compiling Low T-count Oracle Circuits*. [arXiv:1908.01609].
- [30] Yunseong Nam, Neil J. Ross, Yuan Su, Andrew M. Childs & Dmitri Maslov: *“optimiser” Github*. <https://github.com/njross/optimizer>.
- [31] Yunseong Nam, Neil J. Ross, Yuan Su, Andrew M. Childs & Dmitri Maslov (2018): *Automated optimization of large quantum circuits with continuous parameters*. *npj Quantum Information* 4(1), p. 23, doi:10.1038/s41534-018-0072-4. Available at <https://doi.org/10.1038/s41534-018-0072-4>. [arXiv:1710.07345].
- [32] Michael A Nielsen & Isaac Chuang (2000): *Quantum computation and quantum information*. Cambridge University Press, Cambridge UK.
- [33] Renaud Vilmart (2019): *A Near-Optimal Axiomatisation of ZX-Calculus for Pure Qubit Quantum Mechanics*. In: *34th Annual ACM/IEEE Symposium on Logic in Computer Science*, pp. 1–10, doi:10.1109/LICS.2019.8785765. [arXiv:1812.09114].
- [34] Fang Zhang & Jianxin Chen (2019): *Optimizing T gates in Clifford+T circuit as  $\pi/4$  rotations around Paulis*. [arXiv:1903.12456].

## A ZX diagram reference

The ZX calculus — first developed by Coecke and Duncan [15] (see also Refs. [16, 19, 33] for updated treatments, and Refs. [8, 9, 17, 26] for applications to quantum technology) — is a relatively recently developed notation for quantum operations, equipped with rules (the “calculus” part) to compute with this notation. This article does not make explicit use of the “calculus” part of the ZX calculus: while it *does* make statements about equivalences of diagrams which *could* be shown using the calculus, these can and should be understood in the same way that other papers make statements of equivalences of conventional circuit diagrams.

We use ZX notation at various points to describe quantum circuits, including circuits with classically controlled operations and non-local unitaries such as  $\pi/4$ -parity-phase operations. The ZX diagrams in this article can be read merely as a slightly unusual (but convenient) circuit notation. In this Appendix, we provide a reference for this notation, serving also as a glossary of sorts for various operations as they are represented in ZX diagrams, to allow readers to understand our results as well as conventional circuit diagrams would.

### A.1 General statements

For the purposes of this article (and essentially all other practical purposes), ZX diagrams are representations of tensor networks. To represent quantum circuits, it is common to choose a direction in which to read the diagrams from “input” to “output”. (In our paper, we draw these diagrams with input on the left and output on the right, as with the usual circuit notation.) The ZX diagrams of our work are composed of three different kinds of tensor nodes:

- **“Green” nodes** (which are lighter coloured in our article), which may have any number of indices, and as a tensor represents a sort of GHZ state over the standard basis. If a phase parameter  $\theta$  is provided, the tensor also involves a relative phase of  $e^{i\theta}$  between the two terms; otherwise  $\theta = 0$  is assumed (and there is no relative phase).

$$\theta \left( \begin{array}{c} \circ \\ \vdots \\ \circ \end{array} \right) \Bigg\}^n = |0\rangle^{\otimes n} + e^{i\theta} |1\rangle^{\otimes n} \quad (22)$$

In principle, we also permit the border case of  $n = 0$ , in which case this represents the “tensor”  $|0\rangle^{\otimes 0} + e^{i\theta} |1\rangle^{\otimes 0} = (1 + \cos(\theta)) + i \sin(\theta)$ ; though we don’t make use of such nodes in our results.

- **“Red” nodes** (which are darker coloured in our article), which may have any number of indices, and are similar to green nodes except that they are expressed in terms of the  $\{|+\rangle, |-\rangle\}$  basis.

$$\theta \left( \begin{array}{c} \circ \\ \vdots \\ \circ \end{array} \right) \Bigg\}^n = |+\rangle^{\otimes n} + e^{i\theta} |-\rangle^{\otimes n} \quad (23)$$

- **“Hadamard” boxes**, which represent the usual  $2 \times 2$  unitary Hadamard matrix.

$$\boxed{H} = |+\rangle\langle 0| + |-\rangle\langle 1| \quad (24)$$

To represent operations taking some qubits as input, we change of some of the “kets” in the tensor nodes to “bras” — but as  $|0\rangle$ ,  $|1\rangle$ ,  $|+\rangle$  and  $|-\rangle$  are real vectors, this change does not affect any of the tensor coefficients. This allows us to be flexible with our diagrams, and avoid committing to the indices of each node as being explicitly an “input” or an “output”, unless it is a free index of the whole diagram. (In particular, this allows us to draw some closed indices by *vertical* wires, without confusion.)

In the rest of this appendix, we describe some simple examples (and simple extensions) of this notation, which the interested reader should find themselves able to verify by routine calculation.

## A.2 Single-node diagrams

With (light) green or (dark) red nodes of degree 1, we may easily represent states of the  $\{|0\rangle, |1\rangle\}$  basis or  $\{|+\rangle, |-\rangle\}$  basis, albeit supernormalised by a factor of  $\sqrt{2}$ .

$$\bullet\text{---} = |+\rangle^{\otimes 1} + |-\rangle^{\otimes 1} = \sqrt{2}|0\rangle; \quad \pi \bullet\text{---} = |+\rangle^{\otimes 1} - |-\rangle^{\otimes 1} = \sqrt{2}|1\rangle; \quad (25)$$

$$\circ\text{---} = |0\rangle^{\otimes 1} + |1\rangle^{\otimes 1} = \sqrt{2}|+\rangle; \quad \pi \circ\text{---} = |0\rangle^{\otimes 1} - |1\rangle^{\otimes 1} = \sqrt{2}|-\rangle. \quad (26)$$

More generally, green degree-1 nodes may be used to represent newly prepared qubits in the XY plane of the Bloch sphere, and red degree-1 nodes may be used to represent newly prepared qubits in the YZ plane of the Bloch sphere, up to the same supernormalisation of  $\sqrt{2}$ . This additional factor of  $\sqrt{2}$  does not affect our results: the additional factor may be accounted for any time we represent the preparation of a qubit in one of these states.

We may also represent single-qubit measurements by degree-1 nodes oriented in the opposite direction. As re-orienting edges from the right of a node to the left corresponds to turning  $|0\rangle$  to  $\langle 0|$ , turning  $|1\rangle$  to  $\langle 1|$ , and so forth, we then have

$$\text{---}\bullet = \sqrt{2}\langle 0|; \quad \text{---}\bullet\pi = \sqrt{2}\langle 1|; \quad (27)$$

$$\text{---}\circ = \sqrt{2}\langle +|; \quad \text{---}\circ\pi = \sqrt{2}\langle -|. \quad (28)$$

Again, the additional factor of  $\sqrt{2}$  may be accounted for any time we represent a projection of a qubit in one of these states. To represent a measurement which may yield either  $|0\rangle$  or  $|1\rangle$ , or either  $|+\rangle$  or  $|-\rangle$ , we may introduce a variable  $s \in \{0, 1\}$  representing whether a relative phase of  $\pi$  is absent in the result ( $s = 0$ , for the states  $|0\rangle$  or  $|+\rangle$ ) or present in the result ( $s = 1$ , for the states  $|1\rangle$  or  $|-\rangle$ ). We then represent measurement in the  $\{|0\rangle, |1\rangle\}$  basis and the  $\{|+\rangle, |-\rangle\}$  basis respectively as

$$\text{---}\bullet\pi, \{s\} = \langle +| + e^{is\pi}\langle -| \in \{\sqrt{2}\langle 0|, \sqrt{2}\langle 1|\}; \quad (29)$$

$$\text{---}\circ\pi, \{s\} = \langle 0| + e^{is\pi}\langle 1| \in \{\sqrt{2}\langle +|, \sqrt{2}\langle -|\}. \quad (30)$$

The bit  $s$  is in effect a random variable representing the measurement outcome.

In other ZX diagrams (including on nodes of degree 2 or higher), we may use a set  $S = \{s_1, s_2, \dots\}$  in place of the set  $\{s\}$ . This indicates a node in which the presence or absence of the phase of  $\pi$  depends on the parity ( $s_1 \oplus s_2 \oplus \dots$ ) of the entire set  $S$ , rather than on the single bit  $s$ . For example, we may represent Z rotations and X rotations each by a single node of degree 2:

$$\text{---}\overset{\theta}{\circ}\text{---} = |0\rangle\langle 0| + e^{i\theta}|1\rangle\langle 1| = R_z(\theta), \quad \text{---}\overset{\theta}{\bullet}\text{---} = |+\rangle\langle +| + e^{i\theta}|-\rangle\langle -| = R_x(\theta); \quad (31)$$

Then, the following diagrams represent the same operations, conditioned on the parity  $\mathbf{s} = \bigoplus_j s_j$  of a set of bits  $S = \{s_1, s_2, \dots\}$ :

$$\begin{array}{c} \theta, S \\ \text{---}\circ\text{---} \end{array} = R_z(\mathbf{s}\theta) = R_z(\theta)^{\mathbf{s}}, \quad \begin{array}{c} \theta, S \\ \text{---}\bullet\text{---} \end{array} = R_x(\mathbf{s}\theta) = R_x(\theta)^{\mathbf{s}}. \quad (32)$$

This feature of the ZX calculus does not play a prominent role in our work, but is present in our treatment of the Hadamard gadget (Eqn. (15) on page 9) and in principle useful to represent the circuits which we would obtain by representing conditionally-controlled Clifford operations in the ZX calculus.

### A.3 Two-node diagrams

Diagrams of more than one node can be easily constructed simply by composing nodes on their edges. In many cases, this has the same meaning as in conventional quantum circuit diagrams (with the same “feature” that the algebra is read right-to-left, even though the diagram is read left-to-right): for example,

$$\begin{array}{c} \theta \\ \text{---}\circ\text{---} \end{array} \boxed{H} \text{---} = HR_z(\theta) = R_x(\theta)H = \text{---}\boxed{H}\text{---} \begin{array}{c} \theta \\ \bullet \end{array} \quad (33)$$

$$\begin{array}{c} \theta \quad \pi \\ \text{---}\circ\text{---}\bullet\text{---} \end{array} = XR_z(\theta) = R_z(-\theta)X = \text{---}\bullet\text{---}\circ\text{---} \begin{array}{c} \pi \quad -\theta \end{array} \quad (34)$$

As with circuit diagrams, we may also represent the tensor product of operations by representing operations happening on different wires in parallel — for example:

$$\begin{array}{c} \theta \\ \text{---}\circ\text{---} \\ \text{---}\bullet\text{---} \\ \varphi \end{array} = R_z(\theta) \otimes R_x(\varphi). \quad (35)$$

Not all “compositions” of nodes take these forms, however: in general we may compose any two nodes simply by connecting their edges (corresponding to contracting the shared indices of the tensor nodes). An especially important case in point is the way that CNOT operators are represented as ZX terms. As with single-qubit states, the usual representation of CNOT by ZX diagrams is not precisely normalised:


$$\begin{aligned} \begin{array}{c} \text{---}\circ\text{---} \\ \text{---}\bullet\text{---} \end{array} &= |0\rangle\langle 0| \otimes \langle 0|+\rangle \otimes |+\rangle\langle +| + |0\rangle\langle 0| \otimes \langle 0|-\rangle \otimes |-\rangle\langle -| \\ &\quad + |1\rangle\langle 1| \otimes \langle 1|+\rangle \otimes |+\rangle\langle +| + |1\rangle\langle 1| \otimes \langle 1|-\rangle \otimes |-\rangle\langle -| \\ &= \frac{1}{\sqrt{2}}|0\rangle\langle 0| \otimes \mathbb{1} + \frac{1}{\sqrt{2}}|1\rangle\langle 1| \otimes X = \frac{1}{\sqrt{2}}\text{CNOT}. \end{aligned} \quad (36)$$

(Again, the subnormalisation of this diagram does not affect our analysis, and can in principle be accounted for in the ZX representation of any circuit involving CNOT gates.) Note that the shared wire between the red and green dot does not have a specific interpretation as an “input” or an “output” of either — nor is this necessary to provide the interpretation of the diagram as an operator.

### A.4 Multi-node diagrams

Composing the diagrams above, in series or in parallel (and with appropriate accounting for normalisation), suffices to represent an arbitrary unitary operation by the (slightly redundant) gate set consisting of

arbitrary X and Z rotations, Hadamard gates, and CNOT operations. We may also more directly represent somewhat more “exotic” operators using ZX diagrams, and  $\pi/4$ -parity-phase operations are in this case the most relevant example: for instance,



$$\begin{aligned}
&= \left( \langle 0|_d + e^{i\theta} \langle 1|_d \right) \left( |+\rangle_d \langle +++|_{abc} + |-\rangle_d \langle ---|_{abc} \right) \\
&\quad \times \left( |0\rangle\langle 0|_1 \otimes |0\rangle\langle 0|_a + |1\rangle\langle 1|_1 \otimes |1\rangle\langle 1|_a \right) \left( |0\rangle\langle 0|_3 \otimes |0\rangle\langle 0|_b + |1\rangle\langle 1|_3 \otimes |1\rangle\langle 1|_b \right) \\
&\quad \times \left( |0\rangle\langle 0|_5 \otimes |0\rangle\langle 0|_c + |1\rangle\langle 1|_5 \otimes |1\rangle\langle 1|_c \right) \\
&= \frac{1}{2\sqrt{2}} \sum_{a,b,c \in \{0,1\}} \left( \langle 0|_d + e^{i\theta} \langle 1|_d \right) \left( |+\rangle_d + (-1)^{a+b+c} |-\rangle_d \right) \otimes |a,b,c\rangle\langle a,b,c|_{1,3,5} \\
&= \frac{1}{4} \sum_{a,b,c \in \{0,1\}} \left( 1 + e^{i\theta} + (-1)^{a+b+c} - (-1)^{a+b+c} e^{i\theta} \right) |a,b,c\rangle\langle a,b,c|_{1,3,5} \\
&= \frac{1}{2} \sum_{\substack{a,b,c \in \{0,1\} \\ a \oplus b \oplus c = 0}} |a,b,c\rangle\langle a,b,c|_{1,3,5} + \frac{1}{2} \sum_{\substack{a,b,c \in \{0,1\} \\ a \oplus b \oplus c = 1}} e^{i\theta} |a,b,c\rangle\langle a,b,c|_{1,3,5} \\
&= \frac{1}{2} e^{i\theta/2} \exp\left(\frac{1}{2}i\theta(Z \otimes \mathbb{1} \otimes Z \otimes \mathbb{1} \otimes Z)\right).
\end{aligned} \tag{37a}$$

$$\tag{37b}$$

Again, the subnormalisation by a factor of  $\frac{1}{2}$  does not affect our analysis, which is in principle about products of  $D_{S,3}$  operators — merely denoted in our work by these phase gadgets, for convenience — which are proportional to the identity by a global phase.

The existence of rules for transforming ZX diagrams allows us to reason (*i.e.*, to compute) effectively about these diagrams without the need to expand their meaning algebraically as we have been doing in this Appendix. This has particularly motivated our use of the ZX calculus in our work, as a convenient notational tool and also as a means by which we performed our analysis.

For more information about the ZX calculus, and in particular for resources to learn about these diagrammatic computational methods, the interested reader is invited to visit [zxcalculus.com].

## B Details of the `moveH` subroutine and CNOT-commutation heuristic

In this Appendix, we describe our procedures for  $H$  gate extraction and CNOT gate extraction (used in Steps 2 and 4 of the procedure described in Section 4.2) on a high level. For more details, the interested reader may view our source code on Github [<https://github.com/onestruggler/fast-stomp>].

### B.1 The `moveH` subroutine

Our procedure for extracting  $H$  gates from a circuit are built on a subroutine `moveH`, which attempts to move each  $H$  gate as far to the right (the end of the circuit) as possible.

Representing the circuit as a list of gates in a particular order (without parallelisation), this procedure looks for the first  $H$  gate, and attempts to move it to the right. In doing so, it makes use of several simple commutation relations or opportunities for cancellation, for example:

$$\begin{array}{ccc}
 \boxed{H}\boxed{H} \rightarrow \text{---}; & \boxed{H}\boxed{X} \rightarrow \boxed{Z}\boxed{H}; & \boxed{H}\boxed{Z} \rightarrow \boxed{X}\boxed{H}; \\
 \boxed{H}\oplus \rightarrow \text{---}; & \boxed{H} \bullet \rightarrow \oplus \boxed{H}; & \boxed{H} \times \rightarrow \times \boxed{H}.
 \end{array} \tag{38}$$

If the procedure moves the  $H$  gate to a point that it precedes a second  $H$  gate, it proceeds recursively to attempt to move the second  $H$  gate before continuing with the first. When the procedure is finished attempting (successfully or otherwise) to move the second  $H$  gate, it returns to the task of moving the first — moving this gate past the other  $H$  gate, if the attempt to move it ended in failure. This process continues until the procedure has stopped trying to move what originally was the first  $H$  gate.

In attempting to move  $H$  gates, `moveH` may encounter situations in which no progress is possible, without trying to move or cancel other kinds of gates. For instance: in a circuit consisting only of an  $H$  gate followed by four  $T$  gates on a single wire, it is possible to move the  $H$  gate to the end, but only after “pushing” the  $T$  gate which follows it to the right, accumulating the other phase gates to form a  $Z$  gate. In general, if `moveH` encounters a gate  $G$  for which there is no commutation rule provided, it attempts instead to push  $G$  forward, to commute with, accumulate with, or cancel against gates further to its right. In doing so, `moveH` may encounter yet another gate  $F$  for which  $G$  has no provided commutation relation, in which case `moveH` will attempt to move  $F$  further to the right, and so on.

In some cases, there are fruitful opportunities for multi-gate substitutions which either reduce the number of  $H$  gates, or allows an  $H$  gate to be moved further to the right. For instance:

- If in moving an  $H$  gate to the right we encounter an  $S$  gate followed by an  $H$  gate, `moveH` first tries to move the second  $H$  gate. If this fails, we may apply the transformation

$$\boxed{H}\boxed{S}\boxed{H} \rightarrow \boxed{S}\boxed{Z}\boxed{H}\boxed{S}\boxed{Z}. \tag{39}$$

This reduces the number of  $H$  gates by 1. We then move the  $S$  and  $Z$  gates further to the right, then continue by moving the new  $H$  gate to the right.

- If in moving an  $H$  gate to the right we encounter the *control* qubit of a CNOT gate, followed by an  $H$  gate on either the control or target, we again first try to move the  $H$  gate. If this fails, we may apply one of the transformations

$$\begin{array}{ccc}
 \oplus \text{---} & \rightarrow & \boxed{H} \bullet \boxed{H}; \\
 \boxed{H} \bullet \boxed{H} & & \\
 \text{---} & & \\
 \oplus \text{---} & & \\
 \boxed{H} \bullet \boxed{H} & \rightarrow & \boxed{H} \bullet \oplus \text{---}.
 \end{array} \tag{40}$$

This doesn't directly reduce the number of  $H$  gates, but may make it possible to move the later  $H$  gate and the CNOT gate to the right before continuing further, thereby providing an alternative for at least one of the two  $H$  gates to be moved further to the right.

The details of all commutation relations which we define for all of the gates are not important, except that it is important to define these rules in such a way that the procedure terminates (rather than repeatedly commute two gates such as  $T$  and  $CCZ$  past one another, in an attempt to cancel them so that an  $H$  gate can be moved to the right of both). Different techniques will lead to different performances in the ability of `moveH` to reduce the number of  $H$  gates which precede any non-Clifford gate.

## B.2 CNOT movement heuristic

In Step 4, we move all operations which are not single-qubit phase operations or phase-parity operations out of the main body of the circuit. The way that CNOT gates are treated aims, roughly, to avoid generating phase-parity operations on very large subsystems, but does so in a way that attempts to avoid performing too much computation.

The heuristic used to determine which direction to move a CNOT operation is as follows. For each CNOT gate, from the first in the circuit to the last, we compute the following:

1. Compute the set  $P_L$  of all phase-parity gadgets to the left which act on the target but not the control of the CNOT, and the set  $M_L$  of phase-parity gadgets to the left which act on its target and control both.
2. Similarly, compute the set  $P_R$  of phase-parity gadgets to the right which act on the target but not the control of the CNOT, and the set  $M_R$  of phase-parity gadgets to the right which act on its target and control both.
3. If  $P_L - M_L < P_R - M_R$ , we prefer to move the CNOT to the left; otherwise we prefer to move it to the right.

If no other CNOT gate acted on any qubits in common with this left-most CNOT gate, the quantity  $P_L$  (respectively,  $M_L$ ) would correctly indicate how many phase-parity gadgets would act on one more qubit (respectively, one fewer) if we commuted that CNOT to the left. The difference  $P_L - M_L$  then indicates the net change in the cumulative number of qubits acted on by the phase-gadgets to the left of the CNOT. Similar remarks apply for  $P_R - M_R$ , albeit with the important caveat that this figure may be inaccurate if there are further CNOT gates to the right whose targets coincide with the control of the CNOT under consideration.

The approach taken to produce our results is as follows. For the left-most CNOT in the circuit, compute  $P_L$ ,  $M_L$ ,  $P_R$ , and  $M_R$ . Commute the CNOT gate to the left if  $P_L - M_L < P_R - M_R$ , and otherwise to commute it to the right. If in commuting it to the right we encounter another CNOT gate with which it does not commute, we also commute that CNOT gate to the right (and any CNOT gates with which *those* do not commute, *etc.*) Having done this, we compute  $P_L$ ,  $M_L$ ,  $P_R$ , and  $M_R$  for the leftmost remaining CNOT gate in the circuit, where these may depend on the commutations which occurred for the previous CNOT gate. We proceed in this way, recursively from left to right, until no more CNOT gates are in the main body of the circuit.