

ON THE CNOT-COMPLEXITY OF CNOT-PHASE CIRCUITS

MATTHEW AMY^{1,2}, PARSIAZ AZIMZADEH² AND MICHELE MOSCA^{1,3,4,5}

¹ *Institute for Quantum Computing, University of Waterloo, Canada*

² *David R. Cheriton School of Computer Science, University of Waterloo, Canada*

³ *Department of Combinatorics & Optimization, University of Waterloo, Canada*

⁴ *Perimeter Institute for Theoretical Physics, Waterloo, Canada*

⁵ *Canadian Institute for Advanced Research, Toronto, Canada*

ABSTRACT. We study the problem of CNOT-optimal quantum circuit synthesis over gate sets consisting of CNOT and Z -basis rotations of arbitrary angles. We show that the circuit-polynomial correspondence relates such circuits to Fourier expansions of pseudo-Boolean functions, and that for certain classes of functions this expansion uniquely determines the minimum CNOT cost of an implementation. As a corollary we prove that CNOT minimization over CNOT and phase gates is at least as hard as synthesizing a CNOT-optimal circuit computing a set of parities of its inputs. We then show that this problem is NP-complete for two restricted cases where all CNOT gates are required to have the same target, and where the circuit inputs are encoded in a larger state space. The latter case has applications to CNOT optimization over more general Clifford+ T circuits.

We further present an efficient heuristic algorithm for synthesizing circuits over CNOT and Z -basis rotations with small CNOT cost. Our experiments show a 23% reduction of CNOT gates on average across a suite of Clifford+ T benchmark circuits, with a maximum reduction of 43%.

1. INTRODUCTION

The two-qubit controlled-NOT (CNOT) gate forms the backbone of most discrete quantum circuits. As the only entangling operation – and, moreover, the only two-qubit operation – in most common gate sets, it is used judiciously in effectively any practically useful quantum circuit. Even in physical implementations where gates commonly have tunable parameters many use the CNOT gate, or a CNOT gate up to single qubit rotations, as the two qubit entangling gate (e.g. [11]).

Due to the low cost of the controlled-NOT gate in most fault tolerant models [13, 14] compared to single-qubit gates implemented with resource states – for instance, the T gate in the surface code or the Hadamard gate in the 15-qubit Reed-Muller code – reducing the number of CNOT gates is typically regarded as a secondary optimization objective. In some cases, optimizations aimed at reducing T -count or T -depth (the number of layers of parallel T gates in a circuit) result in a massive increase in CNOT count [4]. While in many contexts the savings obtained by optimizing T gates or other such gates makes up for the explosion in CNOT gates, it is nevertheless desirable to find optimization methods which can mitigate this increase, as even in fault tolerant models CNOT gates incur non-negligible cost.

Direct optimization of CNOT gates has generally been studied only in the restricted case of *reversible* circuits. Synthesis and optimization methods have been developed for the subset of *linear* reversible functions, which are exactly the circuits implementable with just CNOT gates [18, 24], but the vast majority of optimizations apply to the NCT gate set, consisting over NOT, CNOT and Toffoli gates, or its generalization to multiply-controlled Toffoli gates. In the former case, additional reductions may be possible between distinct linear reversible circuits separated by other quantum gates. Likewise, further reductions are typically possible between Toffoli or multiply-controlled Toffoli gates once they have been expanded into suitable quantum gate sets [1].

In this work, we consider the problem of minimizing CNOT count in the presence of other quantum gates, specifically phase gates (single qubit Z -basis rotations). By the *circuit-polynomial correspondence* [6, 10, 20], such circuits are known to correspond to weighted sums of parity functions called *phase polynomials* [4].

Using this correspondence, we introduce the problem of synthesizing a *parity network*, a CNOT circuit which computes a set of parities, possibly non-simultaneously. Computing a minimal parity network for a particular set of parity functions is shown to be equivalent to finding a CNOT-optimal circuit *for a particular phase polynomial*, and to be computationally *easier* than the full CNOT optimization problem over CNOT and phase gates. This problem is then shown to be NP-complete in two restricted cases: when all CNOT gates are restricted to the same target bit, and when the m circuit inputs are encoded in the state space of $n > m$ qubits. The former case provides evidence for the hardness of computing minimal parity networks, while the latter case is useful when optimizing CNOT counts in Clifford+T circuits.

We further devise a new heuristic optimization algorithm for CNOT-phase circuits by synthesizing parity networks. The optimization algorithm is inspired by *Gray codes* [15], which cycle through the set of n -bit strings using the exact minimal number of bit flips. Like Gray codes, our algorithm achieves the minimal number of CNOT gates when all 2^n parities are needed. To test our algorithm, we implemented it as a replacement for the T -parallelization sub-routine in the T -par framework [4] to apply our optimization over general Clifford+ T circuits – our experiments show a reduction in CNOT count of 23% on average.

The paper is organized as follows. Section 2 introduces the circuit-polynomial correspondence for quantum circuits composed of CNOT and phase gates, as well as the minimal parity network synthesis problem, then further shows that the general problem of CNOT-minimization over CNOT and phase gates is at least as hard as the parity network synthesis problem. Section 3 studies the complexity of this synthesis problem and shows that the fixed-target and encoded input cases are NP-complete. In Section 4 we present our heuristic CNOT minimization algorithm, and Section 5 gives experimental results.

1.1. Preliminaries. We give a brief overview of the basics of quantum circuits and our notation throughout.

In this paper we work in the circuit model of quantum computation [22]. The state of an n -qubit quantum system is taken as a unit vector in a dimension 2^n complex vector space. As is standard we denote the 2^n basis vectors of the *computational* basis by $|\mathbf{x}\rangle$ for bit strings $\mathbf{x} \in \mathbb{F}_2^n$ – these are called the *classical* states. A general quantum state may then be written as a *superposition* of classical states

$$|\psi\rangle = \sum_{\mathbf{x} \in \mathbb{F}_2^n} \alpha_{\mathbf{x}} |\mathbf{x}\rangle,$$

for complex $\alpha_{\mathbf{x}}$ and having unit norm. The states of two n and m qubit quantum systems $|\psi\rangle$ and $|\phi\rangle$ may be combined into an $n + m$ qubit state by taking their tensor product $|\psi\rangle \otimes |\phi\rangle$ – in particular, $|\mathbf{x}\rangle \otimes |\mathbf{y}\rangle = |\mathbf{x} \parallel \mathbf{y}\rangle$ where \parallel denotes the concatenation of two binary vectors or strings. If to the contrary the state of two qubits cannot be written as a tensor product the two qubits are said to be *entangled*. We write $|\mathbf{x}|$ to denote the *Hamming weight* of the binary string \mathbf{x} .

Quantum circuits, in analogy to classical circuits, carry qubits from left to right along *wires* through *gates* which transform the state. In the unitary circuit model gates are required to implement unitary operators on the state space – that is, quantum gates are modelled by complex-valued matrices U satisfying $UU^\dagger = U^\dagger U = I$, where U^\dagger is the complex conjugate of U . As a result, such quantum computations must be *reversible*, and in particular, the set of classical functions computable by a quantum circuit is the set of invertible, n bit to n bit functions. To perform general (irreversible) classical functions, extra bits called *ancillae* are typically needed to implement them reversibly. We use U_C to denote the unitary operator corresponding to the circuit C and $C :: C'$ to denote the circuit obtained by appending C' to the end of C .

We will primarily be interested in two quantum gates in this paper, the controlled-NOT gate (CNOT) which (as a function of classical states) inverts its second argument conditioned on the value of its first, and the Z -basis rotation $R_Z(\theta)$ which applies a *phase shift* of $e^{i\theta}$ conditioned on its argument. Specifically, we define

$$\text{CNOT}|x\rangle|y\rangle = |x\rangle|x \oplus y\rangle, \quad R_Z(\theta)|x\rangle = e^{2\pi i \theta x} |x\rangle,$$

for $x, y \in \mathbb{F}_2$. By the linearity of quantum computation a unitary operator is fully defined by its effect on the computational basis states. Other common quantum gates include the NOT gate $X|x\rangle = |1 \oplus x\rangle$ and the Hadamard gate $H|x\rangle = \frac{1}{\sqrt{2}} \sum_{y \in \mathbb{F}_2} (-1)^{xy} |y\rangle$. The CNOT gate together with the Hadamard gate and $S = R_Z(\frac{1}{4})$ generate the *Clifford* group, a set of quantum operations particularly important in quantum error correction and for which there are efficient implementations in many fault-tolerant schemes. While the Clifford group alone is not *universal* for quantum computing, in the sense that not every unitary matrix

can be approximated to arbitrary accuracy by a Clifford group circuit, adding the $T = R_Z\left(\frac{1}{8}\right)$ gate to the Clifford group gives a universal gate set. The Clifford group combined with the T gate is typically referred to as the *Clifford+T* gate set. We say a circuit is written over a particular set of gates if the circuit only uses gates from that set.

1.2. Related work. Previous work regarding CNOT optimization has largely focussed on strictly reversible circuits. Iwama, Kambayashi and Yamashita [18] gave some transformation rules which they use to normalize and optimize CNOT-based circuits. More specific to CNOT circuits, Patel, Markov and Hayes [24] gave an algorithm for synthesizing linear reversible circuits which produces circuits of asymptotically optimal size. Their method modifies Gaussian elimination by prioritizing rows which are close in Hamming distance and gives circuits of size at most $O(n^2/\log n)$, coinciding with the known lower bound of $\Theta(n^2/\log n)$ on the worst-case size of CNOT circuits [27].

In the realm of pure quantum circuits, Shende and Markov [26] studied the CNOT cost of Toffoli gates. They proved that 6 CNOT gates is minimal for the Toffoli gate, and gave a lower bound of $2n$ CNOT gates for the n qubit Toffoli gate. More recently, Welch, Greenbaum, Mostame and Aspuru-Guzik [28] studied the construction of efficient circuits for diagonal unitaries. They used similar insights to the ones we use here, notably the use of the 2^n Walsh functions as a basis for n -qubit diagonal operators which correspond to the parity functions in the *Fourier expansions* [23] we use to describe CNOT-phase circuits. While their main objective was to optimize circuits by constructing approximations of the operator which use fewer Walsh functions, they give a construction of an optimal circuit computing all 2^n Walsh functions. They further give CNOT identities which they use to optimize circuits when not all Walsh functions are used, but give no experimental data as to the effectiveness of these optimizations. In contrast, we present and test an algorithm which directly synthesizes an efficient circuit for a specific set of Walsh functions, rather than construct a circuit for the full set and optimize later.

2. CNOT-PHASE CIRCUITS AND FOURIER EXPANSIONS

We begin by introducing the *circuit-polynomial* correspondence [20] for quantum circuits composed of CNOT and R_Z gates, which associates a *phase polynomial* and a linear Boolean transformation with every circuit. While the $\{\text{CNOT}, R_Z\}$ gate set does not contain any branching gates, we call this the *sum-over-paths* form of a circuit in keeping with similar analyses [10, 19, 20].

Definition 2.1. The sum-over-paths form of a circuit C over CNOT and R_Z gates with associated unitary U_C is given by a pseudo-Boolean function

$$f(\mathbf{x}) = \sum_{\mathbf{y} \in \mathbb{F}_2^n} \widehat{f}(\mathbf{y}) \cdot (x_1 y_1 \oplus x_2 y_2 \oplus \cdots \oplus x_n y_n)$$

with coefficients $\widehat{f}(\mathbf{y}) \in \mathbb{R}$, together with a basis state transformation $A \in \text{GL}(n, \mathbb{F}_2)$ such that

$$U_C = \sum_{\mathbf{x} \in \mathbb{F}_2^n} e^{2\pi i f(\mathbf{x})} |A\mathbf{x}\rangle \langle \mathbf{x}|.$$

In particular, U_C maps the basis state $|\mathbf{x}\rangle$ to $e^{2\pi i f(\mathbf{x})} |A\mathbf{x}\rangle$. We refer to the sum-over-paths form as the pair (f, A) .

The expression of $f(\mathbf{x})$ as a weighted sum of parities of \mathbf{x} in Definition 2.1 is the *Fourier expansion*¹ of f with *Fourier coefficients* $\widehat{f}(\mathbf{y})$ [23]. We call the collective set of Fourier coefficients the *Fourier spectrum* or just the spectrum of f . Previous works [4, 6, 8] have used the term *phase polynomial* to refer to the expression of $f(\mathbf{x})$ over the parity basis – we use the term Fourier expansion to avoid confusion with regular polynomial representations of f .

For convenience we define $\chi_{\mathbf{y}} : \mathbb{F}_2^n \rightarrow \mathbb{F}_2$ to be the parity function for the indicator vector $\mathbf{y} \in \mathbb{F}_2^n$. In particular,

$$\chi_{\mathbf{y}}(\mathbf{x}) = x_1 y_1 \oplus x_2 y_2 \oplus \cdots \oplus x_n y_n.$$

¹Standard literature (e.g. [23]) on Fourier analysis of pseudo-Boolean functions uses the multiplicative group $\{-1, 1\}$ as the Boolean group, resulting in a different set of coefficients. As the additive group $\{0, 1\}$ is more natural for quantum computing and the resulting expansion obeys the properties needed, we use this representation. A full discussion can be found in Appendix A.




FIGURE 1. An annotated circuit implementing the doubly controlled Z gate.

The Fourier expansion of f is then written as

$$f(\mathbf{x}) = \sum_{\mathbf{y} \in \mathbb{F}_2^n} \hat{f}(\mathbf{y}) \chi_{\mathbf{y}}(\mathbf{x}).$$

We also define the *support* of the Fourier expansion of f to be the set of parities with non-zero coefficients – that is,

$$\text{supp}(\hat{f}) = \{\mathbf{y} \in \mathbb{F}_2^n \mid \hat{f}(\mathbf{y}) \neq 0\}.$$

It was previously shown by Amy, Maslov and Mosca [4] that there exists a canonical sum-over-paths expression for every $\{\text{CNOT}, T\}$ circuit, and in particular it can be computed in time polynomial in the number of qubits. The same method may be used to compute a canonical sum-over-paths form for any circuit over $\{\text{CNOT}, R_Z\}$, with the only difference being that the Fourier coefficients are defined over \mathbb{R} rather than \mathbb{Z}_8 .

Proposition 2.2. *Every circuit C written over $\{\text{CNOT}, R_Z\}$ has a canonical, polynomial-time computable sum-over-paths form.*

Given a $\{\text{CNOT}, R_Z\}$ circuit, its sum-over-paths form can be computed by first constructing the *annotated circuit*, where the inputs are labelled by x_1, x_2, \dots, x_n , and outputs of every gate are labelled by a parity of the inputs. Note that by the linearity of quantum circuits and the fact that both CNOT and R_Z gates map basis states to basis states (possibly with a phase), the state of each qubit can in fact be described at any point in the circuit as a parity over the input (basis) state. As the CNOT gate maps $|x\rangle|y\rangle$ to $|x\rangle|x \oplus y\rangle$, the output for the control bit of a CNOT gate has the same label as its input, while the target bit is labelled with the XOR of the input labels. An $R_Z(\theta)$ gate does not change the basis state of the qubit, and hence the annotation is unchanged.

Then to construct the sum-over-paths form, for every $R_Z(\theta)$ gate with incoming label $\chi_{\mathbf{y}}(\mathbf{x})$ a factor of $\theta \cdot \chi_{\mathbf{y}}(\mathbf{x})$ is added to $f(\mathbf{x})$ – equivalently, $\hat{f}(\mathbf{y})$ is taken to be the sum of the parameters (θ) of all R_Z gates with incoming label $\chi_{\mathbf{y}}(\mathbf{x})$. The linear transformation A is defined by the mapping $\mathbf{x} \mapsto \mathbf{x}'$ where \mathbf{x}' is the labels at the end of the circuit.

As an example, Fig. 1 shows a circuit implementing the doubly controlled Z operator

$$\Lambda_2(Z) = \sum_{\mathbf{x} \in \mathbb{F}_2^3} e^{2\pi i \frac{x_1 x_2 x_3}{2}} |\mathbf{x}\rangle\langle \mathbf{x}|.$$

The state of each qubit as a parity of an input basis state $|x_1 x_2 x_3\rangle$, $x_1, x_2, x_3 \in \mathbb{F}_2$, is annotated after each gate. Annotations are only shown in cases where the state of the qubit is changed for clarity. Summing up the phase rotations due to R_Z gates gives the expression

$$f(\mathbf{x}) = \frac{1}{8} (x_1 + x_2 + 3(x_1 \oplus x_3) + 3(x_2 \oplus x_3) + (x_1 \oplus x_2 \oplus x_3) + 3(x_1 \oplus x_2) + x_3).$$

As the circuit returns each qubit to its initial state, the sum-over-paths form is (f, I) where $I \in \text{GL}(3, \mathbb{F}_2)$ is the identity matrix.


2.1. Parity networks. A key observation about the canonical sum-over-paths forms is that *only parities which appear in the annotated circuit may have non-zero Fourier coefficient*. Otherwise, the parities may appear in any order, parities may appear in the circuit but not in the sum-over-paths form, or the same parity may appear multiple times. Multiple R_Z gates may also be applied with the same incoming parity throughout a circuit, in which case the Fourier coefficient is the sum of all the rotation angles and can be

replaced with a single R_Z gate – this effect was previously used to optimize T -count in Clifford+ T circuits by merging T gates applied to the same parity [4].

The inverse of the above observation is that a circuit in which every parity in $\text{supp}(\hat{f})$ appears as an annotation can be modified to implement the phase rotation $f(\mathbf{x}) = \sum_{\mathbf{y} \in \mathbb{F}_2^n} \hat{f}(\mathbf{y}) \chi_{\mathbf{y}}(\mathbf{x})$ at no extra CNOT cost. For example, the circuit in Fig. 1 can be modified to give a new circuit with (non-equivalent) sum over paths form (f', I) for

$$f'(\mathbf{x}) = \frac{2}{3}(x_2 \oplus x_3) + \frac{1}{3}(x_1 \oplus x_2 \oplus x_3)$$

with no additional CNOT gates simply by changing the parameters of the fourth and fifth R_Z gates to $\frac{2}{3}$ and $\frac{1}{3}$, respectively, and removing all other phase gates. The resulting circuit is shown below:



This motivates the definition of a *parity network* below as a CNOT circuit computing a set of parities, which can be used to implement phase rotations with Fourier expansions having support contained in that set.

Definition 2.3. A *parity network* for a set $S \subseteq \mathbb{F}_2^n$ is an n -qubit circuit C over CNOT gates where, for each $\mathbf{y} \in S$, the parity $\chi_{\mathbf{y}}(\mathbf{x})$ appears in the annotated circuit.

As all parity networks apply some overall linear transformation of the input, we say a parity network is *pointed at* $A \in \text{GL}(n, \mathbb{F}_2)$ if the overall transformation is A , i.e.,

$$U_C = \sum_{\mathbf{x} \in \mathbb{F}_2^n} |A\mathbf{x}\rangle \langle \mathbf{x}|.$$

For convenience we refer to a parity network with the trivial transformation as an *identity parity network*. In the context of synthesizing parity networks, we use the term *pointed parity network* to refer to a parity network applying a specific linear transformation.

We can now formalize the above observations with the following proposition, stating that the problem of finding a minimal size (pointed) parity network is equivalent to finding a CNOT-minimal circuit having a particular sum-over-paths form. For the remainder of the paper we consider these problems interchangeable.

Proposition 2.4. Given a circuit C over $\{\text{CNOT}, R_Z\}$ with sum-over-paths form (f, A) , the circuit C' obtained from C by removing all phase gates is a parity network for $\text{supp}(\hat{f})$ pointed at A .

Furthermore, given parity network C for $S \subseteq \mathbb{F}_2^n$ pointed at A , the circuit C' obtained from C by, for every $\mathbf{y} \in S$, inserting a phase gate $R_Z(\hat{f}(\mathbf{y}))$ where $\chi_{\mathbf{y}}(\mathbf{x})$ appears as an annotation, has sum-over-paths form (f, A) .

Proof. For the former statement, by definition of the canonical sum-over-paths form the annotated version of C necessarily contains $\chi_{\mathbf{y}}(\mathbf{x})$ as a label for every $\mathbf{y} \in \text{supp}(\hat{f})$. Since R_Z gates do not change labels or permute basis vectors, the circuit C' contains all the labels of C and implements the permutation A , hence C' is a parity network $\text{supp}(\hat{f})$ pointed at A .

The proof of the latter statement is similar. □

2.2. From CNOT-minimization to parity network synthesis. Proposition 2.4 implies that the problem of finding a minimal pointed size parity network is equivalent to the problem of finding a CNOT-minimal circuit for a particular sum-over-paths form. However, it is not necessarily the case that a CNOT-minimal circuit having a particular sum-over-paths form is a CNOT-minimal circuit implementing a particular unitary matrix. Since for any integer-valued function $k : \mathbb{F}_2^n \rightarrow \mathbb{Z}$,

$$e^{2\pi i f(\mathbf{x})} = e^{2\pi i (f(\mathbf{x}) + k(\mathbf{x}))},$$

it may in general be possible to instead synthesize a *different* sum-over-paths form giving the same unitary operator, but with lower CNOT cost. For instance,

$$\frac{1}{2}(x_1 \oplus x_2), \quad \text{and} \quad \frac{1}{2}x_1 + \frac{1}{2}x_2$$




FIGURE 2. Relationship between circuit, sum-over-paths and unitary representations. As an example, the circuits C, C' and C'' all correspond to the same unitary operator U , while only C and C' have the same sum-over-paths representation, (f, A) . In either case, the basis state transformation A is the same.

differ by an integer-valued function, $k(x_1, x_2) = x_1x_2$, and hence implement the same phase rotation. The left expression, together with the identity basis state transformation, gives rise to a minimal circuit containing 2 CNOT gates, while the expression on the right requires no CNOT gates to implement, at the expense of an extra phase gate. The two annotated circuits are shown below.




Figure 2 summarizes the relationship between circuits, sum-over-paths, and unitaries. As we are concerned with the question of minimizing CNOT gates over circuits with equal *unitary* representations, a natural question is how this relates to the question of minimizing CNOT gates over circuits with equal *sum-over-paths* representations. We now show that so long as no rotation gates have angles which are *dyadic fractions* – numbers of the form $\frac{a}{2^b}$ where a and b are integers – the problems coincide.

We first formalize the intuition that two sum-over-paths forms correspond to equivalent unitaries if and only if their phases are related by an integer-valued function. In particular, we define the equivalence class of a phase function as

$$[f] = \{f' : \mathbb{F}_2^n \rightarrow \mathbb{R} \mid f' = f + k \text{ where } k : \mathbb{F}_2^n \rightarrow \mathbb{Z}\},$$

and we say f' is equivalent to f , written $f' \sim f$, if $f' \in [f]$. From these definitions the following proposition follows straightforwardly.

Proposition 2.5. *Given $f, f' : \mathbb{F}_2^n \rightarrow \mathbb{R}$ and $A, A' \in \text{GL}(n, \mathbb{F}_2)$, the unitary matrices*

$$\sum_{\mathbf{x} \in \mathbb{F}_2^n} e^{2\pi i f(\mathbf{x})} |A\mathbf{x}\rangle \langle \mathbf{x}|, \quad \sum_{\mathbf{x} \in \mathbb{F}_2^n} e^{2\pi i f'(\mathbf{x})} |A'\mathbf{x}\rangle \langle \mathbf{x}|$$

are equal if and only if $A' = A$ and $f' \sim f$.

It is a known fact² [23] that every pseudo-Boolean function $f : \mathbb{F}_2^n \rightarrow \mathbb{R}$ has a unique Fourier expansion. However, to study the relationship between Fourier expansions of equivalent functions, it will be important to know their precise form in the case of integer-valued functions.

Proposition 2.6. *For any integer-valued function $k : \mathbb{F}_2^n \rightarrow \mathbb{Z}$, the Fourier coefficients of k are dyadic fractions.*

Proof. Let $k : \mathbb{F}_2^n \rightarrow \mathbb{Z}$ be an integer-valued pseudo-Boolean function. It is known [16] that k has a unique representation as an n -ary multilinear polynomial over \mathbb{Z} – that is,

$$k(\mathbf{x}) = \sum_{\mathbf{y} \in \mathbb{F}_2^n} a_{\mathbf{y}} \mathbf{x}^{\mathbf{y}}$$

where $\mathbf{x}^{\mathbf{y}} = x_1^{y_1} x_2^{y_2} \cdots x_n^{y_n}$ and $a_{\mathbf{y}} \in \mathbb{Z}$ for all \mathbf{y} .

Using the identity $x + y - (x \oplus y) = 2xy$ for $x, y \in \mathbb{F}_2$, we can derive an inclusion-exclusion formula for the monomial $\mathbf{x}^{\mathbf{y}}$ which we prove explicitly in Appendix B:

$$(1) \quad 2^{|\mathbf{y}|-1} \mathbf{x}^{\mathbf{y}} = \sum_{\mathbf{y}' \subseteq \mathbf{y}} (-1)^{|\mathbf{y}'|-1} \chi_{\mathbf{y}'}(\mathbf{x}).$$

Note that binary vectors are viewed as subsets of $\{1, \dots, n\}$ for convenience. Since $a_{\mathbf{y}} \in \mathbb{Z}$ for all \mathbf{y} and dyadic fractions are closed under addition, we observe that the Fourier coefficients of $k(\mathbf{x})$ are dyadic fractions:

$$k(\mathbf{x}) = \sum_{\mathbf{y} \in \mathbb{F}_2^n} a_{\mathbf{y}} \mathbf{x}^{\mathbf{y}} = \sum_{\mathbf{y} \in \mathbb{F}_2^n} \left(\sum_{\mathbf{y}' \subseteq \mathbf{y}} (-1)^{|\mathbf{y}'|-1} \frac{a_{\mathbf{y}}}{2^{|\mathbf{y}|-1}} \right) \chi_{\mathbf{y}'}(\mathbf{x}).$$

□

The proof of Proposition 2.6 also suffices to prove a more general result, namely that any function from \mathbb{F}_2^n to an Abelian group G in which 2 is a regular element has a unique Fourier expansion over G . This more general version also subsumes a similar result proven in [6], namely that Fourier expansions are unique over any cyclic group of order co-prime to 2; indeed, 2 is regular in any such group.

We next use the above proposition to show that any pseudo-Boolean function with non-dyadic spectrum has a property of *minimal support* over all equivalent functions. This is important as a parity network for S is also a parity network for any subset of S .

Proposition 2.7. *Let $f : \mathbb{F}_2^n \rightarrow \mathbb{R}$ be a pseudo-Boolean function having a Fourier spectrum not containing any non-zero dyadic fractions. Then for any $f' \sim f$,*

$$\text{supp}(\widehat{f}) \subseteq \text{supp}(\widehat{f}').$$

Proof. Consider some pseudo-Boolean function f' such that $f' \sim f$. By definition we have $f' = f + k$ for some function $k : \mathbb{F}_2^n \rightarrow \mathbb{Z}$. Expanding $f(\mathbf{x})$ and $k(\mathbf{x})$ with their Fourier expansions we have

$$f'(\mathbf{x}) = f(\mathbf{x}) + k(\mathbf{x}) = \sum_{\mathbf{y} \in \mathbb{F}_2^n} (\widehat{f}(\mathbf{y}) + \widehat{k}(\mathbf{y})) \chi_{\mathbf{y}}(\mathbf{x}).$$

Now since for any \mathbf{y} , $\widehat{k}(\mathbf{y}) = \frac{a_{\mathbf{y}}}{2^{|\mathbf{y}|}}$, $\widehat{f}(\mathbf{y}) + \widehat{k}(\mathbf{y}) \neq 0$. Thus $\text{supp}(\widehat{f}) \subseteq \text{supp}(\widehat{f}')$ as required. □

We can now prove that the problem of synthesizing a minimal (pointed) parity network is at least as hard as general CNOT-minimization. As a corollary, synthesizing a minimal parity network solves the CNOT-minimization problem whenever the rotation angles, and hence the Fourier coefficients, are not dyadic fractions.

Theorem 2.8. *Given $A \in \text{GL}(n, \mathbb{F}_2)$, the problem of finding a minimal parity network for $S \subseteq \mathbb{F}_2^n$ pointed at A reduces (in polynomial time) to the problem of finding a CNOT-minimal circuit equivalent to an n -qubit circuit C over $\{\text{CNOT}, R_Z\}$.*

²As we use a slightly different definition of Fourier expansion, not every pseudo-Boolean function has a unique Fourier expansion. Our Fourier expansions are only unique up to constant terms, which correspond to global phase factors and are not directly synthesizable over $\{\text{CNOT}, R_Z\}$.

Proof. Given $A \in \text{GL}(n, \mathbb{F}_2)$ and $S \subseteq \mathbb{F}_2^n$, define $f : \mathbb{F}_2^n \rightarrow \mathbb{R}$ as

$$f(\mathbf{x}) = \sum_{\mathbf{y} \in S} \frac{1}{3} \chi_{\mathbf{y}}(\mathbf{x}).$$

It is known [4] that a circuit C over $\{\text{CNOT}, R_Z\}$ implementing the sum-over-paths form (f, A) can be constructed in polynomial time.

Now let C' be a CNOT-minimal circuit equivalent to C . By Proposition 2.5 the sum-over-paths form of C' must be (f', A) for some $f' \sim f$. However, by Proposition 2.7, $S = \text{supp}(\hat{f}) \subseteq \text{supp}(\hat{f}')$, so by definition, the circuit obtained from C' by removing all R_Z gates is a (necessarily minimal) parity network for S pointed at A . \square

Remark 2.9. In cases when the Fourier coefficients contain dyadic fractions, it may in general be possible to further minimize CNOT-count by optimizing over all equivalent phase functions. This question was studied in [6] from the perspective of T -count optimization – in that case, the number of T gates depends only on the size of the support, $|\text{supp}(\hat{f})|$. By contrast, the size of the support of a Fourier spectrum does not necessarily correspond to the size of a minimal parity network – e.g., $\frac{1}{2}(x_1 \oplus x_2)$ has smaller support than $\frac{1}{2}x_1 + \frac{1}{2}x_2$ but a larger minimal parity network as shown earlier – which appears to make the problem of minimizing CNOTs size over all equivalent functions more difficult.

3. COMPLEXITY OF PARITY NETWORK MINIMIZATION

We now turn to the question of the complexity of finding minimal parity networks. We study two cases in particular where the problem can be shown to be NP-complete – the fixed-target case, and with encoded inputs (i.e. with ancillae). At the end of the section we discuss the case of synthesizing a minimal parity network with arbitrary targets and no ancillae.

Note that we focus on the problem of synthesizing minimal parity networks rather than *pointed* parity networks – that is, synthesizing a minimal parity network up to some arbitrary overall linear transformation. However, in the cases we consider equivalent reductions for pointed parity networks are also possible.

3.1. Fixed-target minimal parity network. We call the problem of synthesizing a minimal parity network in which every CNOT gate has the same target the *fixed-target minimal parity network problem*. Formally, we define the associated decision problem MPNP_{FT} below:

- PROBLEM: Fixed-target minimal parity network (MPNP_{FT})
 INSTANCE: A set of strings $S \subseteq \mathbb{F}_2^n$, and a positive integer k .
 QUESTION: Does there exist an n -qubit circuit C over CNOT gates of length at most k such that C is a parity network for S ?

Remark 3.1. In general not every set of strings S admits an ancilla-free fixed-target parity network, as the value of the target bit necessarily appears in every parity calculation of a fixed-target CNOT circuit. It follows that an (ancilla-free) parity network for S is synthesizable if and only if there exists an index i such that for every $\mathbf{y} \in S$, $y_i = 1$. However, a fixed-target parity network may always be synthesized by adding a single ancillary bit initialized to the state $|0\rangle$. In particular, given a set $S \subseteq \mathbb{F}_2^n$ and $A \in \text{GL}(n, \mathbb{F}_2)$, we may construct

$$S' = \{(\mathbf{y} \parallel 1) \mid \mathbf{y} \in S\},$$

where $\mathbf{y} \parallel 1$ denotes the length $n + 1$ string obtained by concatenating \mathbf{y} with 1. It may then be observed that a fixed-target parity network for S' is always synthesizable, and in particular forms a parity network for S when the $(n + 1)$ th bit is initialized to $|0\rangle$.

To show that the fixed-target minimal parity network problem is NP-complete, we introduce the *Hamming salesman problem* (HTSP) [12]. Recall that the n -dimensional hypercube is the graph with vertices $\mathbf{x} \in \mathbb{F}_2^n$ and edges between $\mathbf{x}, \mathbf{y} \in \mathbb{F}_2^n$ if \mathbf{x} and \mathbf{y} differ in one coordinate (i.e. have Hamming distance 1).

PROBLEM: Hamming salesman (HTSP)
INSTANCE: A set of strings $S \subseteq \mathbb{F}_2^n$, and a positive integer k .
QUESTION: Does there exist a path in the n -dimensional hypercube of length at most k starting at $\mathbf{0}$ and going through each vertex $\mathbf{y} \in S$?

An equivalent (from a complexity theoretic viewpoint) version of the Hamming salesman problem exists where a cycle rather than path is found. Intuitively, the Hamming salesman problem is to find a sequence of at most k bit-flips iterating through every string in some set S starting from the initial string $00\dots 0$. In the case when $S = \mathbb{F}_2^n$ the minimal number of bit flips is known to be 2^n , corresponding to one bit flip per string; this is the well known *Gray code*, a total ordering on \mathbb{F}_2^n where each subsequent string differs by exactly one bit. We will come back to this connection later in Section 4 when designing a synthesis algorithm.

Ernvall, Katajainen and Penttonen [12] show that the Hamming salesman problem is in fact NP-complete, hence we can use a reduction from HTSP to prove NP-completeness of MPNP_{FT} .

Theorem 3.2. *MPNP_{FT} is NP-complete.*

Proof. Clearly MPNP_{FT} is in NP, as the state of each bit as a parity of the input values at each state in a CNOT circuit is polynomial-time computable [4], and hence a parity network can be efficiently verified. Since HTSP is NP-complete [12] it then suffices to show NP-hardness by reducing the Hamming salesman problem to the fixed-target minimal parity network problem.

Given an instance $(S \subseteq \mathbb{F}_2^n, k)$ of HTSP, we construct an instance $(S' \subseteq \mathbb{F}_2^{n+1}, k')$ of MPNP_{FT} with size polynomial in $|S| \cdot n$ as follows:

$$S' = \{(\mathbf{x} \parallel 1) \mid \mathbf{x} \in S\}, \quad k' = k.$$

Suppose there exists a fixed-target parity network C for S' with length at most k . Without loss of generality we may assume that the fixed target is the $(n+1)$ th bit, as if some $i \neq n+1$ is the target index, then for all $\mathbf{y} \in S'$, $y_i = 1 = y_{n+1}$ and hence swapping bits i and $n+1$ yields a parity network for S' . We can then construct a length k hypercube path through each vertex $\mathbf{y} \in S$ with starting point $\mathbf{0}$ by mapping C to a sequence of bit flips on each CNOT's control bit. Indeed, by noting that

$$\text{CNOT}|x_i\rangle|x_{n+1} \oplus \chi_{\mathbf{y}}(\mathbf{x})\rangle = |x_i\rangle|x_{n+1} \oplus \chi_{\mathbf{y} \oplus \{i\}}(\mathbf{x})\rangle,$$

where $\{i\}$ is taken as the bitstring \mathbf{z} with $z_i = 1, z_{j \neq i} = 0$, each CNOT gate in C with control i has the affect of flipping the i th bit of \mathbf{y} . By the definition of a parity network, for every $\mathbf{y} \in S$, the parity

$$x_{n+1} \oplus \chi_{\mathbf{y}}(\mathbf{x})$$

appears as an annotation in the circuit, in particular on the $(n+1)$ th bit which had initial state $x_{n+1} \oplus \chi_{\mathbf{0}}(\mathbf{x})$, hence the sequence of bit flips passes through each vertex $\mathbf{y} \in S$ starting from $\mathbf{0}$.

Likewise, if there exists a length k tour through each $\mathbf{y} \in S$, given by a sequence of bit flips, the circuit defined by mapping each bit flip on i to a CNOT with control i and target $n+1$ is a length k parity network for S' and A' . \square

As the minimum k for which a parity network exists is at most $(n-1) \cdot |S|$, the optimization version of MPNP_{FT} is also in NP, and hence is NP-complete.

Corollary 3.3. *The problem of finding a minimal fixed-target parity network is NP-complete.*

It may be observed that the proof of Theorem 3.2 can be modified to show that the problem of finding a minimal *pointed* parity network with fixed CNOT targets is also NP-complete. In particular, taking A to be the identity transformation gives a reduction from the cycle version of HTSP.

3.2. Minimal parity network with encoded inputs. As CNOT-phase circuits form a relatively small [3], classically simulable group, one of the main applications of their optimization is to optimize sub-circuits of circuits over more powerful gate sets. The T -par optimization algorithm [4] previously took this approach, re-synthesizing $\{\text{CNOT}, T\}$ subcircuits of Clifford+ T circuits with optimal T -depth. As Selinger [25] showed, adding ancillae increases the amount of T gate parallelization possible when performing this re-synthesis.

In the case of CNOT optimization, the situation is similar in that ancillae can have a significant effect on the number of CNOT gates required to implement a parity network, particularly if the ancillae are initialized in linear combinations of the primary inputs. For instance, the boxed CNOT-phase sub-circuit below on the

left performs a phase rotation of $\frac{1}{8}(x_2 \oplus x_3)$ – by noting that the ancilla begins the sub-circuit already in the state $x_2 \oplus x_3$ we can remove both CNOT gates, as shown by the equivalent circuit on the right.



We now consider the problem of synthesizing minimal parity networks when some of the inputs are linear combinations of others. Formally, given a linear transformation $E \in \mathbb{F}_2^{m \times n}$ where $m > n$, a string $\mathbf{w} \in \mathbb{F}_2^m$ we say is an *encoding* of $\mathbf{x} \in \mathbb{F}_2^n$ if $E\mathbf{x} = \mathbf{w}$. The *minimal parity network with encoded inputs problem* (MPNP_E) is then to find a parity network for a given set $S \subseteq \mathbb{F}_2^m$ as a function of the primary inputs $\mathbf{x} \in \mathbb{F}_2^n$, but with the initial state $|E\mathbf{x}\rangle$ rather than $|\mathbf{x}\rangle$.

- PROBLEM: Minimal parity network with encoded inputs (MPNP_E)
 INSTANCE: A set of strings $S \subseteq \mathbb{F}_2^m$, a linear transformation $E \in \mathbb{F}_2^{m \times n}$, and a positive integer k .
 QUESTION: Does there exist an m -qubit circuit C over CNOT gates of length at most k such that C is a parity network for some set $S' \subseteq \mathbb{F}_2^m$ where for any $\mathbf{y} \in S$ there exists $\mathbf{w} \in S'$ such that $E^T \mathbf{w} = \mathbf{y}$?

It can be observed that a parity network for some set S' as above is equivalent to a parity network for S starting from the initial state $|E\mathbf{x}\rangle$ for any $\mathbf{x} \in \mathbb{F}_2^n$. In particular, for any $\mathbf{w} \in S'$ and $\mathbf{x} \in \mathbb{F}_2^n$,

$$\chi_{\mathbf{w}}(E\mathbf{x}) = \mathbf{w}^T E\mathbf{x} = \chi_{E^T \mathbf{w}}(\mathbf{x}) = \chi_{\mathbf{y}}(\mathbf{x}).$$

MPNP_E is again NP-complete, which we prove by a reduction from the well known NP-complete *Maximum-likelihood Decoding Problem* (MLDP). While the focus of this paper is on $\{\text{CNOT}, R_Z\}$ circuits, the proof may be modified to establish the more general result that synthesizing a CNOT circuit implementing $A \in \mathbb{F}_2^{m \times n}$ is NP-complete.

- PROBLEM: Maximum-likelihood decoding (MLDP)
 INSTANCE: A linear transformation $H \in \mathbb{F}_2^{m \times n}$, a vector $\mathbf{y} \in \mathbb{F}_2^m$, and a positive integer k .
 QUESTION: Does there exist a vector $\mathbf{w} \in \mathbb{F}_2^n$ of weight at most k such that $H\mathbf{w} = \mathbf{y}$?

In the case when H is the parity check matrix of a code C and \mathbf{y} is the syndrome of some vector \mathbf{z} , finding the minimum such \mathbf{w} gives the minimum weight vector in the coset of $\mathbf{z} + C$, corresponding to a minimum distance decoding of \mathbf{z} . Berlekamp, McEliece, and van Tilborg [7] proved that the Maximum-likelihood decoding problem is NP-complete, and so we may reduce it to the minimal parity network with encoded inputs problem to show NP-completeness.

Theorem 3.4. *MPNP_E is NP-complete.*

Proof. As noted in the proof of Theorem 3.2, MPNP_E is clearly in NP since a parity network can be verified in polynomial time. To establish NP-hardness we give a reduction from MLDP.

Given an instance (H, \mathbf{y}, k) of MLDP we construct an instance (S, k') of MPNP_E as follows:

$$S = \{\mathbf{y}\}, \quad E = H^T, \quad k' = k - 1.$$

Suppose there exists a vector $\mathbf{w} \in \mathbb{F}_2^n$ of weight at most k such that $H\mathbf{w} = \mathbf{y}$. Then we know $S' = \{\mathbf{w}\}$ satisfies the requirement that for any $\mathbf{y} \in S$ there exists $\mathbf{w} \in S'$ such that $E^T \mathbf{w} = H\mathbf{w} = \mathbf{y}$. Moreover, the parity computation $\chi_{\mathbf{w}}$ can be computed with $|\mathbf{w}| - 1 \leq k'$ CNOT gates, hence there exists a parity network of length at most k' for S'

On the other hand, suppose there exists a length $\leq k'$ parity network for some set S' where there exists $\mathbf{w} \in S'$ such that $E^T \mathbf{w} = H\mathbf{w} = \mathbf{y}$. By noting that

$$\text{CNOT}|\chi_{\mathbf{y}}(\mathbf{x})\rangle|\chi_{\mathbf{z}}(\mathbf{x})\rangle = |\chi_{\mathbf{y}}(\mathbf{x})\rangle|\chi_{\mathbf{y} \oplus \mathbf{z}}(\mathbf{x})\rangle,$$

$|\mathbf{y} \oplus \mathbf{z}| \leq |\mathbf{y}| + |\mathbf{z}|$. As each bit starts in some state $x_i = \chi_{\{x_i\}}(\mathbf{x})$ we see that the size of the parity in any bit at any point in the parity network is at most $k' + 1 \leq k$, and so $|\mathbf{w}| \leq k$ as required. \square

Corollary 3.5. *The problem of finding a minimal parity network with encoded inputs is NP-hard.*

As in the fixed-target case, the problem of finding a minimal pointed parity network with encoded inputs is also NP-complete, by virtue of the fact that a minimal identity parity network for a singleton set $S = \{\mathbf{y}\}$ necessarily has the form $C :: C'$ where both C and $(C')^{-1}$ are both parity network for S . Recall that the inverse circuit $(C')^{-1}$ has the same length as C' .

3.3. Discussion. While the case of parity network synthesis with encoded inputs corresponds to the practical cases of synthesis with ancillae and sub-circuit re-synthesis, it relies on the hardness of finding a minimal sum of linearly dependent vectors, or minimum distance decoding. It would appear that the problem becomes easier when using unencoded inputs, as each vector and hence parity may be expressed uniquely over the inputs. We leave the complexity of finding a minimal parity network without ancillae as an open problem, but conjecture that the identity version is at least as hard as synthesizing a minimal fixed-target identity parity network. In particular, it appears that the two problems coincide whenever one bit appears in every parity.

4. A HEURISTIC SYNTHESIS ALGORITHM

In this section we present an efficient, heuristic algorithm for synthesizing small parity networks. The algorithm is inspired by Gray codes, which were noted in Section 3 to iterate through all 2^n elements of \mathbb{F}_2^n minimally with one bit flip per string. The situation is different for synthesizing parity networks as the bits which can be flipped depend on the state of all n bits, so our method works by trying to find subsets of S which can be efficiently iterated with a Gray code on a fixed target. In the limit where $S = \mathbb{F}_2^n$, the algorithm gives a minimal size parity network for S . Again we focus just on the problem of synthesizing a parity network up to some arbitrary overall linear transformation.

The algorithm GRAY-SYNTH, is presented in pseudo-code in Algorithm 1. Note that $\text{CNOT}_{i,j}$ denotes a CNOT gate with control i and target j , and $E_{i,j}$ denotes the elementary \mathbb{F}_2 -matrix adding row i to row j . Given a set of binary strings S , the algorithm synthesizes a parity network for S by repeatedly choosing an index i to expand and then effectively recurring on the co-factors S_0 and S_1 , consisting of the strings $\mathbf{y} \in S$ with $y_i = 0$ or 1, respectively. As a subset S is recursively expanded, CNOT gates are applied so that a designated *target* bit contains the (partial) parity $\chi_{\mathbf{y}}(\mathbf{x})$ where $y_i = 1$ if and only if $y'_i = 1$ for all $\mathbf{y}' \in S$ – if S is a singleton $\{\mathbf{y}'\}$, then $\mathbf{y} = \mathbf{y}'$, hence the target bit contains the value $\chi_{\mathbf{y}'}(\mathbf{x})$ as desired. Notably, rather than *uncomputing* this sequence of CNOT gates when a subset S is finished being synthesized, the algorithm maintains the invariant that the remaining parities to be computed are expressed over the current state of the bits. This allows the algorithm to avoid the “backtracking” inherent in uncomputing-based methods.

More precisely, the invariant of Algorithm 1 described above is expressed in the following lemma.

Lemma 4.1. *Let C be a CNOT circuit and $S \subseteq \mathbb{F}_2^n$. For any positive integer i we let $C_{\leq i}$ denote the first i gates of C , c_i and t_i be the control and target of the i th CNOT gate in C . If we define*

$$\begin{aligned} A_0 &= I & \mathbf{y}_0 &= \mathbf{y} \\ A_i &= E_{c_i, t_i} A_{i-1} & \mathbf{y}_i &= E_{t_i, c_i} \mathbf{y}_{i-1} \end{aligned}$$

for every $\mathbf{y} \in S$, then it follows that for any $\mathbf{x} \in \mathbb{F}_2^n$, $U_{C_{\leq i}}|\mathbf{x}\rangle = |A_i \mathbf{x}\rangle$ and

$$\chi_{\mathbf{y}_i}(A_i \mathbf{x}) = \chi_{\mathbf{y}}(\mathbf{x}).$$

Proof. The fact that $U_{C_{\leq i}}|\mathbf{x}\rangle = |A_i \mathbf{x}\rangle$ follows simply from the fact that $\text{CNOT}_{i,j}|\mathbf{x}\rangle = |E_{i,j} \mathbf{x}\rangle$. For the latter fact, clearly $\chi_{\mathbf{y}_i}(A_i \mathbf{x}) = \chi_{\mathbf{y}}(\mathbf{x})$ by definition. Moreover, recall that

$$\chi_{\mathbf{y}}(A \mathbf{x}) = \mathbf{y}^T A \mathbf{x} = \chi_{A^T \mathbf{y}}(\mathbf{x})$$

Algorithm 1 Algorithm for synthesizing a parity network

```

1: function GRAY-SYNTH( $S \subseteq \mathbb{F}_2^n$ )
2:   New empty circuit  $C$ 
3:   New empty stack  $Q$ 
4:    $Q.\text{push}(S, \{1, \dots, n\}, \epsilon)$ 
5:   while  $Q$  non-empty do
6:      $(S, I, i) \leftarrow Q.\text{pop}$ 
7:     if  $S = \emptyset$  or  $I = \emptyset$  then return
8:     else if  $i \in \mathbb{N}$  then
9:       while  $\exists j \neq i \in \{1, \dots, n\}$  s.t.  $\mathbf{y}_j = 1$  for all  $\mathbf{y} \in S$  do
10:         $C \leftarrow C :: \text{CNOT}_{j,i}$ 
11:        for all  $(S', I', i') \in Q \cup (S, I, i)$  do
12:          for all  $\mathbf{y} \in S'$  do
13:             $\mathbf{y} \leftarrow E_{i,j} \mathbf{y}$ 
14:          end for
15:        end for
16:      end while
17:    end if
18:     $j \leftarrow \arg \max_{j \in I} \max_{x \in \mathbb{F}_2} |\{\mathbf{y} \in S \mid y_j = x\}|$ 
19:     $S_0 \leftarrow \{\mathbf{y} \in S \mid y_j = 0\}$ 
20:     $S_1 \leftarrow \{\mathbf{y} \in S \mid y_j = 1\}$ 
21:    if  $i \in \{\epsilon\}$  then
22:       $Q.\text{push}(S_1, I \setminus \{j\}, j)$ 
23:    else
24:       $Q.\text{push}(S_1, I \setminus \{j\}, i)$ 
25:    end if
26:     $Q.\text{push}(S_0, I \setminus \{j\}, i)$ 
27:  end while
28:  return  $C$ 
29: end function

```

and hence by induction,

$$\begin{aligned}
\chi_{\mathbf{y}_i}(A_i \mathbf{x}) &= \chi_{(A_i)^T \mathbf{y}_i}(\mathbf{x}) \\
&= \chi_{(A_{i-1})^T E_{t_i, c_i} E_{t_i, c_i} \mathbf{y}_{i-1}}(\mathbf{x}) \\
&= \chi_{(A_{i-1})^T \mathbf{y}_{i-1}}(\mathbf{x}) \\
&= \chi_{\mathbf{y}}(\mathbf{x}).
\end{aligned}$$

□

It is clear to see that each \mathbf{y}_i for $\mathbf{y} \in S$ in Lemma 4.1 is the value of \mathbf{y} after i iterations. To see then that the output is in fact a parity network for S , it suffices to observe that whenever $S = \{\mathbf{y}_i\}$ and $I = \emptyset$, $|\mathbf{y}_i| = 1$ and thus by Lemma 4.1, some bit is in the state $|\chi_{\mathbf{y}_i}(A_i \mathbf{x})\rangle = |\chi_{\mathbf{y}}(\mathbf{x})\rangle$ after the i th CNOT gate. While the fact that $|\mathbf{y}_i| = 1$ is assured by lines 9-16 in this case, the non-zero elements of the target strings are actually zero-ed out earlier, as the algorithm expands each coordinate. In particular, the first time the “1” branch is taken when expanding a set, corresponding to the first 1 seen over the indices previously examined, the target bit i is set – taking further “1” branches result in the row j being flipped to 0 with a single CNOT. In this way the algorithm makes use of the redundancy in Fourier spectrum S .

In practice, a parity network implementing some particular basis state transformation is typically needed. We take the approach of synthesizing pointed parity network by first synthesizing a regular parity network, then implementing the remaining linear transformation – i.e. AA_i^{-1} where A_i is the linear transformation implemented by the network. In our implementation we use the Patel-Markov-Hayes algorithm [24] which gives asymptotically optimal CNOT count.

While the correctness of Algorithm 1 is independent of the choice of index j to expand in line 18, in practice it has a large impact on the size of the resulting parity network. We chose j so as to maximize the size of the largest subset, S_0 or S_1 , i.e. $j = \arg \max_{j \in I} \max_{x \in \mathbb{F}_2^n} |\{\mathbf{y} \in S \mid y_j = x\}|$. The intuition behind this choice is that as a subset S of \mathbb{F}_2^n with m bits fixed approaches $|S| = 2^{n-m}$, the minimal parity network for S approaches one CNOT per string, corresponding to the Gray code in the limit. We also ran experiments with other methods of choosing j ; we found that $j = \arg \max_{j \in I} \max_{x \in \mathbb{F}_2^n} |\{\mathbf{y} \in S \mid y_j = x\}|$ gave the best results on average.

4.1. Examples.

Example 4.2. To illustrate Algorithm 1, we demonstrate the use of GRAY-SYNTH to synthesize a circuit over $\{\text{CNOT}, R_Z\}$ implementing the diagonal operator $U|\mathbf{x}\rangle = e^{2\pi i f(\mathbf{x})}|\mathbf{x}\rangle$ given by

$$f(\mathbf{x}) = \frac{1}{8} [(x_2 \oplus x_3) + x_1 + (x_1 \oplus x_4) + (x_1 \oplus x_2 \oplus x_3) + (x_1 \oplus x_2 \oplus x_4) + (x_1 \oplus x_2)].$$

$$\begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \end{pmatrix} \quad \begin{array}{l} x_1 - x_1 \\ x_2 - x_2 \\ x_3 - x_3 \\ x_4 - x_4 \end{array}$$

Starting with the initial set S , written as the columns of the matrix on the left, we choose a bit maximizing the number of 0's or 1's in S . As $j = 1$ in this case, we construct the cofactors S_0 and S_1 on the values in the first row, and recurse on S_0 .

$$\begin{pmatrix} \boxed{0} & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \end{pmatrix} \quad \begin{array}{l} x_1 - x_1 \\ x_2 - x_2 \\ x_3 - x_3 \\ x_4 - x_4 \end{array}$$

The greyed row above indicates that the first row has been partitioned, and the box indicates the current subset being synthesized. The algorithm next selects row 2 and immediately descends into the 1-cofactor, since $S_0 = \emptyset$. Again, the algorithm selects row 3, and since both rows 2 and 3 have the value 1, a CNOT is applied with bit 2 as the target and 3 as the control. The remaining vectors are updated by multiplying with $E_{2,3}$ – the modified entries are shown in red.

$$\begin{pmatrix} \boxed{0} & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \end{pmatrix} \quad \begin{array}{l} x_1 - x_1 \\ x_2 - x_2 \\ x_3 - x_3 \\ x_4 - x_4 \end{array}$$


$$\begin{pmatrix} \boxed{0} & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 1 \\ \boxed{1} & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \end{pmatrix} \rightarrow \begin{pmatrix} \boxed{0} & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 1 \\ \boxed{0} & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 \end{pmatrix} \quad \begin{array}{l} x_1 - x_1 \\ x_2 \oplus x_3 \\ x_3 - x_3 \\ x_4 - x_4 \end{array}$$

As the final row has the value 0, we're finished with this column and may continue with the remaining subsets. The single 1 in row 2 indicates that the second qubit currently holds the value of the corresponding parity $x_2 \oplus x_3$, as seen in the circuit on the right.


$$\begin{pmatrix} \boxed{0} & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 \end{pmatrix} \quad \begin{array}{l} x_1 - x_1 \\ x_2 \oplus x_3 \\ x_3 - x_3 \\ x_4 - x_4 \end{array}$$

Again the algorithm chooses row 2 maximizing the number of entries which are the same for the remaining columns, and recurses on the 0-cofactor as shown below.

$$\begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 \end{pmatrix}$$




$$\begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 \end{pmatrix}$$




In expanding the last row, we first examine the 0-cofactor and find nothing to do, then the 1-cofactor, at which point we need to apply a CNOT with target bit 1 and control 4.

$$\begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 \end{pmatrix}$$




$$\begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 \end{pmatrix}$$




Now backtracking and entering the 1-cofactor for the remaining columns, we find we need to apply a CNOT between bits 2 and 1 to zero out the row.

$$\begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 \end{pmatrix}$$




Continuing on we recurse on the 0-cofactor of row 3 and apply a CNOT with target bit 1, control 4, before backtracking to the 1-cofactor.

$$\begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 \end{pmatrix}$$




$$\begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}$$




For the remaining two columns we first zero out row 3 by applying a CNOT gate between bits 3 and 1, then finally descend into the cofactors on the last row.

$$\begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix} \rightarrow \begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}$$



$$\begin{pmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}$$






FIGURE 3. Circuit implementing the doubly-controlled Z gate $\Lambda_2(Z)$ synthesized with Algorithm 1. The CNOT-minimal Fourier expansion in this case gives $S = \mathbb{F}_2^3 \setminus \{000\}$, and $A = I$.





FIGURE 4. Annotated parity network for the set $S = \{(\mathbf{y} \parallel 1) \mid \mathbf{y} \in \mathbb{F}_2^3\}$. Note that the parity network corresponds exactly to the Gray code on \mathbb{F}_2^3 .


$$\left(\begin{array}{cccccc} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{array} \right) \rightarrow \left(\begin{array}{cccccc} 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{array} \right)$$

The overall linear transformation applied is

$$A = \begin{bmatrix} 1 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$



so the algorithm completes by appending a circuit computing A^{-1} . Inserting $T = R_Z(1/8)$ gates in the relevant positions, we get the following circuit computing $|\mathbf{x}\rangle \mapsto e^{2\pi i f(\mathbf{x})} |\mathbf{x}\rangle$:



Example 4.3. Figure 3 shows a circuit implementing the doubly-controlled Z gate, corresponding to an identity parity network for $\mathbb{F}_2^3 \setminus \{000\}$, synthesized with Algorithm 1 followed by the Patel-Markov-Hayes algorithm. In this case both the parity network and the identity parity network are minimal, as verified by brute force search – further, the circuit synthesized by Algorithm 1 reproduces exactly the minimal circuit for $\mathbb{F}_2^3 \setminus \{000\}$ from [28]. In general, for any n the identity parity network for $\mathbb{F}_2^n \setminus \{\mathbf{0}\}$ synthesized in this manner has the same structure, using $2^n - 2$ CNOT gates, compared to 2^n bit flips for the Gray code.

Example 4.4. If instead of $S = \mathbb{F}_2^n \setminus \{\mathbf{0}\}$ we have $S \simeq \mathbb{F}_2^m$ for some $m < n$, Algorithm 1 instead gives a circuit corresponding directly to the Gray code. In particular, Fig. 4 shows a parity network for $S = \{(\mathbf{y} \parallel 1) \mid \mathbf{y} \in \mathbb{F}_2^3\} \simeq \mathbb{F}_2^3$ synthesized with Algorithm 1, where $S \simeq \mathbb{F}_2^3$. In this case it can be observed that the controls of the CNOT gates are exactly the bits flipped in a Gray code for \mathbb{F}_2^3 . Further, it may be noted that this is a minimal size parity network for S , and is in fact a fixed-target parity network.

4.2. Synthesis with encoded inputs. Given an encoder $E \in \mathbb{F}_2^{m \times n}$, we can use Algorithm 1 to synthesize a parity network for some set S with encoded inputs as follows. Recall that a parity network for $S \subseteq \mathbb{F}_2^n$ with inputs encoded by E corresponds to a parity network for some set $S' \subseteq \mathbb{F}_2^m$ such that for any $\mathbf{y} \in S$,




FIGURE 5. Average CNOT counts of parity networks as computed by Algorithm 1 and brute force minimization over all sets of 4 bit parities.

there exists $\mathbf{w} \in S'$ where $E^T \mathbf{w} = \mathbf{y}$. While finding the minimal such \mathbf{w} would require solving the NP-hard Maximum-likelihood decoding problem, we can efficiently compute *some* \mathbf{w} using a *generalized inverse*.

Recall that a generalized inverse A^g of a matrix $A \in \mathbb{F}_2^{m \times n}$ is an n by m matrix over \mathbb{F}_2 such that

$$AA^gA = A,$$

and in particular $AA^g\mathbf{y} = \mathbf{y}$ whenever there exists \mathbf{x} such that $A\mathbf{x} = \mathbf{y}$. We use such a generalized inverse rather than a specific inverse such as the Moore-Penrose pseudoinverse, as the latter typically does not exist over finite fields. By contrast, for $A \in \mathbb{F}_2^{m \times n}$, a generalized inverse A^g may be computed [9] by finding invertible matrices P, Q over \mathbb{F}_2 such that

$$A = P \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix} Q$$

and computing

$$A^g = Q^{-1} \begin{bmatrix} I_r & 0 \\ 0 & 0 \end{bmatrix} P^{-1}.$$

P and Q may likewise be found by first reducing A to row-echelon form, then reducing its transpose to row-echelon form.

It is worth noting that it may be possible to perform additional optimization by optimizing the set S' with A^g . In particular, it is known [9] that the set of solutions to the linear system $A\mathbf{x} = \mathbf{y}$ is given by $\{A^g\mathbf{y} + (I - A^gA)\mathbf{w} \mid \mathbf{w} \in \mathbb{F}_2^m\}$, which may be possible to optimize with classical techniques. We tried brute-force optimizing the set S' for some small instances and found negligible effects on overall CNOT counts, though we leave it as an open question as to whether scalable sub-optimal methods reduce parity network sizes in large benchmarks.

5. EVALUATION

We implemented Algorithm 1 in Haskell in the open-source quantum circuit toolkit FEYNMAN³. Experiments were run in Debian Linux running on a quad-core 64-bit Intel Core i7 2.40 GHz processor and 8 GB RAM.

We generated all 4-bit minimal parity networks by brute force search and compared them with the parity networks generated with Algorithm 1. Figure 5 graphs the results, with CNOT cost averaged over sets S of parities with the same size. The results show that our algorithm synthesizes optimal or near-optimal networks for small and large sets S , and diverges slightly for sets of parities containing around half of the possible parities. The divergence peaks at $|S| = 8$, exactly half of the 2^4 parities, with Algorithm 1 coming within 15% of the minimal number of CNOT gates on average. On examining the structure of optimal parity networks for sets on which Algorithm 1 performed poorly, it appears that the optimal parity networks save on CNOT cost by making more judicious use of *shortcuts* – leaving qubits in particular states to flip between distant parities in other bits quickly. In general we found that the optimal results in these cases *were not* achievable just by

³<https://github.com/meamy/feynman>

using the GRAY-SYNTH algorithm with different index expansion orders. An effective synthesis algorithm may then be to combine Algorithm 1 for small and large sets with a different heuristic for sets S of size close to 2^{n-1} .

5.1. Benchmarks. To evaluate the performance of Algorithm 1 on practical quantum circuits, we implemented a variant of the T -par algorithm [4] with Algorithm 1 replacing the original Matroid partitioning-based $\{\text{CNOT}, R_Z\}$ synthesis, which optimized T -depth rather than CNOT-count. Circuits written over the Clifford+ T gate set are analyzed in sum-over-paths form and effectively split into alternating sequences of $\{\text{CNOT}, R_Z\}$ gates and Hadamard gates, the former of which are synthesized with Algorithm 1. As the input basis state of a particular $\{\text{CNOT}, R_Z\}$ sub-circuit is expressed by an encoding of the entire circuit’s inputs as well as *path variables* [4, 10], we use the method of Section 4.2 to synthesize the circuit. As Algorithm 1 is only used to perform the $\{\text{CNOT}, R_Z\}$ synthesis sub-routine, our optimized circuits have the same T -count as the original T -par algorithm.

We ran our implementation on a suite of benchmarks previously used to evaluate quantum circuit optimizations [4, 6]. Benchmark circuits containing Toffoli gates and multiply controlled Toffoli gates were expanded to the Clifford+ T gate set using the Toffoli gate decomposition from [5] with 7 CNOT gates, and the Nielsen-Chuang multiply controlled Toffoli decomposition [22] giving $2(k-2)$ Toffoli gates per k -controlled Toffoli. While other decompositions, particularly with better CNOT counts exist, we chose these standard decompositions to compare against previous results, which used the same decompositions. All but 4 optimized circuits⁴ (Mod-Adder_1024, Cycle_17_3, GF(2^{64})-Mult and HWB_8) were formally verified to be correct using the method of [2], with verification failing on the remaining 4 due to lack of memory. To the best of the authors’ knowledge, this also represents the first quantum circuit optimization work where a majority of the optimized circuits have also been formally verified.

Table 1 reports the results of our experiments. On average, Algorithm 1 resulted in an 23% reduction of CNOT gates, with 43% reduction in the best case. In reality there may be more CNOT reduction on average, as the algorithm performed relatively poorly on the Galois field multipliers, which comprise over a quarter of the benchmarks. Further, only 4 benchmarks took over a second to complete, lending evidence to the scalability of our method. One benchmark, CSLA-MUX_3, did observe an *increase* in CNOT gates of 31% – this appears to be due to our sub-optimal method of generating a pointed parity network from a parity network, combined with the fact that very few T gates cancel (that is, the support of the Fourier expansions synthesized have total size close to the number of T gates in the original circuit). It may be possible to reduce the overhead in this case, and further reduce CNOT counts in the other benchmarks, by synthesizing pointed parity network directly, rather than synthesizing a parity network followed by a linear permutation.

As Algorithm 1 effectively replaces the matroid partitioning algorithm used in T -par [4], optimizing CNOT count rather than T -depth, we compared the time, CNOT and T -depth tradeoffs of the two methods. The results are presented in Table 1. While Algorithm 1 reduces the CNOT-count significantly compared to matroid partitioning, the T -depth is significantly increased in those cases, as expected. In fact, the T -depth and CNOT count appear to be inversely proportional – however, unlike matroid partitioning which increases the CNOT count of all benchmarks, Algorithm 1 frequently reduces the T -depth of the original circuit. The runtimes are also significantly reduced in large circuits with a large number of T gates, which is expected as the asymptotic complexity of the matroid partitioning algorithm used in T -par is significantly higher [4].

We also compared the results of our heuristic optimization algorithm to a recent heuristic by Nam, Ross, Su, Childs and Maslov [21], which was brought to our attention while preparing this manuscript. While their algorithm does not explicitly look at the problem of CNOT minimization, they achieve significant CNOT reductions by combining (among other techniques) a T -par style phase gate folding stage with optimized decompositions of specific gates and rule-based local rewrites. While their software is not open-source, when available the “light” optimization results reported in [21] are given in Table 1, as that algorithm most closely matches the scalability of ours. Algorithm 1 typically results in similar CNOT counts to their circuit optimizer, with Algorithm 1 reporting better CNOT counts on some circuits (e.g., VE-Adder_3, CSUM-MUX_9) and worse on others (e.g., CSLA-MUX_3, QCLA-Mod_7). Additionally, it may be noted that in the case of the Galois field multipliers, the reductions Nam et al. achieve are from using base circuits with fewer CNOT

⁴It was reported in [2] that 4 separate optimized benchmark circuits (CSLA-MUX_3, Adder_8, Mod-Mult_55 and GF(2^{32})) provably contained errors. These errors have since been fixed and these circuits now pass verification.

Benchmark	n	Base			T -par (matroids)			Nam et al. (L)			T -par (GRAY-SYNTH)					
		CNOT	T	T -depth	Time (s)	CNOT	T	T -depth	Time	CNOT	T	Time	CNOT	T	T -depth	% Red.
Grover_5	9	336	336	144	0.028	425	154	59	–	–	–	0.001	226	154	128	32.7
Mod 5_4	5	32	28	12	0.001	50	16	7	< 0.001	28	16	0.001	26	16	15	18.8
VBE-Adder_3	10	80	70	24	0.004	91	24	9	< 0.001	50	24	0.004	46	24	22	42.5
CSLA-MUX_3	15	90	70	21	0.011	213	62	13	< 0.001	76	64	0.073	118	62	29	-31.1
CSUM-MUX_9	30	196	196	18	0.038	268	84	16	< 0.001	168	84	0.095	148	84	23	24.5
QCLA-Com_7	24	215	203	27	0.055	271	94	20	0.001	132	95	0.097	136	94	28	36.7
QCLA-Mod_7	26	441	413	66	0.099	729	237	58	0.004	302	237	0.145	360	237	67	18.4
QCLA-Adder_10	36	267	238	24	0.076	442	162	19	0.002	195	162	0.112	214	162	30	19.9
Adder_8	24	466	399	69	0.082	654	215	54	0.004	331	215	0.165	359	215	73	23.0
RC-Adder_6	14	104	77	33	0.012	133	47	14	< 0.001	73	47	0.080	71	47	36	31.7
Mod-Red_21	11	122	119	48	0.076	191	73	28	< 0.001	81	73	0.091	86	73	59	29.5
Mod-Mult_55	9	55	49	15	0.003	89	35	9	< 0.001	40	35	0.004	40	35	20	27.3
Mod-Adder_1024	28	2005	1995	831	0.608	2842	1011	249	–	–	–	0.739	1390	1011	863	30.7
Mod-Adder_1048576	58	16680	16660	7292	39.565	26600	7339	5426	–	–	–	12.272	11080	7339	6570	33.6
Cycle 17_3	35	4532	4739	2001	2.431	6547	1955	315	–	–	–	2.618	2968	1955	1857	37.4
GF(2 ⁴)-Mult	12	115	112	36	0.006	197	68	14	0.001	99	68	0.041	106	68	39	7.8
GF(2 ⁵)-Mult	15	179	175	51	0.012	334	111	17	0.001	154	115	0.038	163	111	53	8.9
GF(2 ⁶)-Mult	18	257	252	60	0.019	484	150	22	0.003	221	150	0.055	235	150	63	8.6
GF(2 ⁷)-Mult	21	349	343	72	0.044	713	217	27	0.004	300	217	0.450	319	217	75	8.6
GF(2 ⁸)-Mult	24	469	448	84	0.047	932	264	31	0.006	405	264	0.066	428	264	87	8.7
GF(2 ⁹)-Mult	27	575	567	96	0.113	1198	351	34	0.010	494	351	0.076	526	351	95	8.5
GF(2 ¹⁰)-Mult	30	709	700	108	0.147	1462	410	36	0.009	609	410	0.081	648	410	109	8.6
GF(2 ¹⁶)-Mult	48	1837	1792	180	1.519	4556	1040	65	0.065	1581	1040	0.363	1691	1040	585	7.9
GF(2 ³²)-Mult	96	7292	7168	372	132.133	22205	4128	126	1.834	6299	4128	5.571	6636	4128	2190	9.0
GF(2 ⁶⁴)-Mult	192	28861	28672	756	19072.290	105830	16448	256	58.341	24765	16448	114.310	25934	16448	7716	10.1
Ham_15 (low)	17	259	161	69	0.069	376	97	38	–	–	–	0.043	208	97	83	19.7
Ham_15 (med)	17	616	574	240	0.108	695	242	98	–	–	–	0.089	357	242	201	42.0
Ham_15 (high)	20	2500	2457	996	0.498	3036	1021	424	–	–	–	0.376	1502	1021	837	39.9
HWB_6	7	131	105	45	0.006	199	75	25	–	–	–	0.029	110	75	63	16.0
HWB_8	12	7508	5425	2106	1.879	12428	3531	951	–	–	–	1.706	6861	3531	2752	8.6
QFT_4	5	48	69	48	0.006	76	67	45	–	–	–	0.005	48	67	62	0.0
$\Lambda_3(X)$	5	21	21	9	0.001	32	15	8	< 0.001	14	15	< 0.001	14	15	12	33.3
$\Lambda_3(X)$ (Barenco)	5	28	28	12	0.015	29	16	7	< 0.001	20	16	< 0.001	18	16	16	35.7
$\Lambda_4(X)$	7	35	35	15	0.002	55	23	7	< 0.001	22	23	0.001	22	23	18	37.1
$\Lambda_4(X)$ (Barenco)	7	56	56	24	0.002	71	28	11	< 0.001	40	28	< 0.001	36	28	26	35.7
$\Lambda_5(X)$	9	49	49	21	0.004	76	31	10	< 0.001	30	31	0.003	30	31	24	38.8
$\Lambda_5(X)$ (Barenco)	9	84	84	36	0.004	103	40	15	< 0.001	60	40	< 0.001	54	40	34	35.7
$\Lambda_{10}(X)$	19	119	119	51	0.066	167	71	21	< 0.001	70	71	0.071	70	71	54	41.2
$\Lambda_{10}(X)$ (Barenco)	19	224	224	96	0.024	271	100	25	0.001	160	100	0.029	144	100	81	35.7
Total																22.6

TABLE 1. Benchmark optimization results. Base gives the original circuit statistics, T -par (matroids) gives optimization results using [4], Nam et al. (L) reports the light optimization results from [21] (where available), and T -par (GRAY-SYNTH) gives the results using Algorithm 1 instead of matroid partitioning. The % reduction in CNOT gates over the base circuit is reported in the last column.

gates. As their optimizations rely largely on special-purpose synthesis and local rewrites, the techniques are complementary and so it may be possible to combine both to further reduce CNOT counts. Further, most of their circuits are not verified in any way and hence may have errors, in comparison to ours which are almost all formally certified to be correct.

Nam et al. also report on the results of a “heavy” optimization algorithm which generally performs slightly better than Algorithm 1, with the exception of the VBE-adder_3 and Mod 5_4 benchmarks. This optimization however does not scale to the largest of our benchmark circuits, such as $GF(2^{64})$ -Mult, as a result of the use of local rule-based rewrites to optimize CNOT-phase subcircuits. An interesting question is whether first performing a rough re-synthesis with Algorithm 1 would reduce run-times for the “heavy” local rewrites.

6. CONCLUSION

In this paper we have shown that the problem of synthesizing a CNOT-minimal circuit over CNOT and arbitrary angle R_Z gates, is at least as hard as synthesizing a minimal size CNOT circuit computing a set of parities of the inputs. We have moreover shown that this problem is in fact NP-complete for the cases when all CNOT gates have a fixed target, and when the inputs are encoded in some larger state space. As a result it would appear that the problem of optimizing CNOT counts in quantum circuits is intractable. We further presented a heuristic algorithm to solve this problem which is inspired by Gray codes. The algorithm, when used in a sum-over-paths quantum circuit optimizer (T -par) to re-synthesize CNOT and R_Z sub-circuits, reduces CNOT counts by 23% on average on our set of benchmarks. While the heuristic does not significantly outperform special-purpose synthesis and rule-based rewriting techniques [21] on arithmetic benchmarks, our technique is novel and complementary to rewriting, opening new avenues for optimizing CNOT gates by direct synthesis of parity networks.

While we suspect that the problem of synthesizing minimal parity networks is intractable, we leave the exact complexity as an open question. The two-dimensional nature of the problem makes it an unnatural problem for reductions, as efficient circuits make use of “shortcuts” by computing strategic parities into various bits. On the other hand, it may turn out to be possible to find an efficient algorithm for synthesizing minimal parity networks. Likewise, while we have presented an effective heuristic algorithm for synthesizing parity networks, we leave developing a heuristic algorithm synthesizing small pointed parity networks directly as a possibility for future work.

As a final point of consideration, while we have studied the problem of optimizing CNOT gates in an unrestricted architecture, physical chip designs typically have limited connectivity and can only apply CNOT gates between connected qubits. While arbitrary CNOT gates can be implemented with a sequence of CNOT gates having length at most proportional to the diameter of the connectivity graph, it nonetheless remains possible that a better circuit may be found by synthesizing one directly for a given topology. A natural and important direction for future research is then to find CNOT optimization algorithms which take into account such connectivity constraints, a problem which likewise appears to be computationally intractable [17]. Going in the other direction, it may be possible to find more optimizations at the logical level by taking into account particular fault-tolerant models – for instance, in lattice surgery-based circuits where multi-target CNOT gates have cost strictly less than separate CNOT gates [17].

REFERENCES

- [1] N. Abdessaied, M. Amy, R. Drechsler, and M. Soeken. Complexity of reversible circuits and their quantum implementations. *Theoretical Computer Science*, 618:85–106, 2016.
- [2] M. Amy. Towards Large-Scale Functional Verification of Universal Quantum Circuits. In *Proceedings of the 15th International Conference on Quantum Physics and Logic (QPL'18)*, To appear, 2018, 1805.06908.
- [3] M. Amy, J. Chen, and N. J. Ross. A Finite Presentation of CNOT-Dihedral Operators. In *Proceedings of the 14th International Conference on Quantum Physics and Logic (QPL'17)*, pages 84–97, 2017, 1701.00140.
- [4] M. Amy, D. Maslov, and M. Mosca. Polynomial-time t-depth optimization of Clifford+T circuits via matroid partitioning. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 33(10):1476–1489, Oct 2014, arXiv:1303.2042.
- [5] M. Amy, D. Maslov, M. Mosca, and M. Roetteler. A meet-in-the-middle algorithm for fast synthesis of depth-optimal quantum circuits. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, 32(6):818–830, 2013, arXiv:1206.0758.
- [6] M. Amy and M. Mosca. T-count optimization and reed-muller codes. *arXiv preprint*, 2016, arXiv:1601.07363.

- [7] E. Berlekamp, R. McEliece, and H. van Tilborg. On the inherent intractability of certain coding problems (corresp.). *IEEE Trans. Inf. Theor.*, 24(3):384–386, Sept. 2006.
- [8] E. T. Campbell and M. Howard. Unifying gate synthesis and magic state distillation. *Physical Review Letters*, 118(6):060501, 2017, arXiv:1606.01906.
- [9] S. Campbell and C. Meyer. *Generalized Inverses of Linear Transformations*. Classics in Applied Mathematics. Society for Industrial and Applied Mathematics (SIAM, 3600 Market Street, Floor 6, Philadelphia, PA 19104), 2009.
- [10] C. M. Dawson, A. P. Hines, D. Mortimer, H. L. Haselgrove, M. A. Nielsen, and T. J. Osborne. Quantum computing and polynomial equations over the finite field \mathbb{Z}_2 . *Quantum Info. Comput.*, 5(2):102–112, Mar. 2005, quant-ph/0408129.
- [11] S. Debnath, N. Linke, C. Figgatt, K. Landsman, K. Wright, and C. Monroe. Demonstration of a small programmable quantum computer with atomic qubits. *Nature*, 536(7614):63–66, 2016, arXiv:1603.04512.
- [12] J. Ernvall, J. Katajainen, and M. Penttonen. NP-completeness of the Hamming salesman problem. *BIT Numerical Mathematics*, 25(1):289–292, 1985.
- [13] A. G. Fowler, A. M. Stephens, and P. Groszkowski. High-threshold universal quantum computation on the surface code. *Phys. Rev. A*, 80:052312, Nov 2009, arXiv:0803.0272.
- [14] D. Gottesman. Theory of fault-tolerant quantum computation. *Phys. Rev. A*, 57:127–137, Jan 1998, quant-ph/9702029.
- [15] F. Gray. Pulse code communication, Mar. 17 1953. US Patent 2,632,058.
- [16] P. Hammer and S. Rudeanu. *Boolean methods in operations research and related areas*. Ökonometrie und Unternehmensforschung. Springer-Verlag, 1968.
- [17] D. Herr, F. Nori, and S. J. Devitt. Optimization of lattice surgery is np-hard. *npj Quantum Information*, 3(1):35, 2017, 1702.00591.
- [18] K. Iwama, Y. Kambayashi, and S. Yamashita. Transformation rules for designing cnot-based quantum circuits. In *Proceedings of the 39th Annual Design Automation Conference, DAC '02*, pages 419–424, New York, NY, USA, 2002. ACM.
- [19] D. E. Koh, M. D. Penney, and R. W. Spekkens. Computing quopit clifford circuit amplitudes by the sum-over-paths technique. *Quantum Info. Comput.*, 17(13&14):1081–1095, Nov. 2017, arXiv:1702.03316.
- [20] A. Montanaro. Quantum circuits and low-degree polynomials over \mathbb{F}_2 . *Journal of Physics A: Mathematical and Theoretical*, 50(8):084002, 2017, arXiv:1607.08473.
- [21] Y. Nam, N. J. Ross, Y. Su, A. M. Childs, and D. Maslov. Automated optimization of large quantum circuits with continuous parameters. *npj Quantum Information*, 4(1):23, 2018, arXiv:1710.07345.
- [22] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge Series on Information and the Natural Sciences. Cambridge University Press, 2000.
- [23] R. O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, New York, NY, USA, 2014.
- [24] K. N. Patel, I. L. Markov, and J. P. Hayes. Optimal synthesis of linear reversible circuits. *Quantum Info. Comput.*, 8(3):282–294, Mar. 2008, quant-ph/0302002.
- [25] P. Selinger. Quantum circuits of T -depth one. *Phys. Rev. A*, 87:042302, Apr 2013, arXiv:1210.0974.
- [26] V. V. Shende and I. L. Markov. On the CNOT-cost of TOFFOLI gates. *Quantum Info. Comput.*, 9(5):461–486, May 2009, arXiv:0803.2316.
- [27] V. V. Shende, A. K. Prasad, I. L. Markov, and J. P. Hayes. Reversible logic circuit synthesis. In *Proceedings of the 2002 IEEE/ACM International Conference on Computer-aided Design, ICCAD '02*, pages 353–360, New York, NY, USA, 2002. ACM, quant-ph/0207001.
- [28] J. Welch, D. Greenbaum, S. Mostame, and A. Aspuru-Guzik. Efficient quantum circuits for diagonal unitaries without ancillas. *New Journal of Physics*, 16(3):033040, 2014, arXiv:1306.3991.

APPENDIX A. FOURIER EXPANSIONS

The Fourier expansion of a pseudo-Boolean function used in this paper is not the one typically used in the analysis of Boolean functions [23]. In particular, the Boolean group is typically taken as the multiplicative group $\{-1, 1\}$ with the parity function $\chi_{\mathbf{y}} : \{-1, 1\}^n \rightarrow \{-1, 1\}$ defined by $x_1^{y_1} x_2^{y_2} \cdots x_n^{y_n}$. The resulting expansion is then a *multi-linear* polynomial in \mathbf{x} .

We can recover the standard Fourier expansion $f(\mathbf{x}) = \sum_{\mathbf{y} \in \mathbb{F}_2^n} \tilde{f}(\mathbf{y}) \tilde{\chi}_{\mathbf{y}}(\mathbf{x})$ by defining $\tilde{\chi}_{\mathbf{y}} : \mathbb{F}_2^n \rightarrow \{-1, 1\}$ as

$$\tilde{\chi}_{\mathbf{y}}(\mathbf{x}) = (-1)^{x_1^{y_1} x_2^{y_2} \cdots x_n^{y_n}},$$

in which case by observing that

$$\hat{f}(\mathbf{y}) \chi_{\mathbf{y}}(\mathbf{x}) = \frac{1 + \hat{f}(\mathbf{y})}{2} \tilde{\chi}_{\mathbf{y}}(\mathbf{x})$$

we see that

$$\tilde{f}(\mathbf{0}) = \frac{2^n - 1}{2}, \quad \tilde{f}(\mathbf{y}) = \frac{1}{2} \hat{f}(\mathbf{y}).$$

APPENDIX B. PROOF OF EQ. (1)

In this section we prove Eq. (1), namely that

$$2^{|\mathbf{y}|-1} \mathbf{x}^{\mathbf{y}} = \sum_{\mathbf{y}' \subseteq \mathbf{y}} (-1)^{|\mathbf{y}'|-1} \chi_{\mathbf{y}'}(\mathbf{x})$$

for any $\mathbf{x}, \mathbf{y} \in \mathbb{F}_2^n$.

Lemma B.1. *For any $\mathbf{x} \in \mathbb{F}_2^n$,*

$$2^{n-1} x_1 x_2 \cdots x_n = \sum_{\mathbf{y} \in \mathbb{F}_2^n} (-1)^{|\mathbf{y}|-1} \chi_{\mathbf{y}}(\mathbf{x})$$

Proof. Clearly the formula is satisfied for $n = 1$. Now consider $n = k + 1$ for some k . Using the identity $x + y - (x \oplus y) = 2xy$ for any $x, y \in \mathbb{F}_2$ and basic arithmetic we observe that

$$\begin{aligned} 2^k x_1 x_2 \cdots x_{k+1} &= 2^{k-1} x_1 x_2 \cdots (2x_k x_{k+1}) \\ &= 2^{k-1} x_1 x_2 \cdots (x_k + x_{k+1} - (x_k \oplus x_{k+1})) \\ &= 2^{k-1} x_1 x_2 \cdots x_k + 2^{k-1} x_1 x_2 \cdots x_{k+1} - 2^{k-1} x_1 x_2 \cdots (x_k \oplus x_{k+1}) \end{aligned}$$

Next we define the length k vectors $\mathbf{x}', \mathbf{x}'', \mathbf{x}''' \in \mathbb{F}_2^k$ as follows:

$$x'_i = x_i, \quad x''_i = \begin{cases} x_{k+1} & \text{if } i = k \\ x_i & \text{otherwise} \end{cases}, \quad x'''_i = \begin{cases} x_k \oplus x_{k+1} & \text{if } i = k \\ x_i & \text{otherwise} \end{cases}$$

By induction we see that

$$\begin{aligned} &2^k x_1 x_2 \cdots x_{k+1} \\ &= 2^{k-1} x'_1 x'_2 \cdots x'_k + 2^{k-1} x''_1 x''_2 \cdots x''_k - 2^{k-1} x'''_1 x'''_2 \cdots x'''_k \\ &= \sum_{\mathbf{y} \in \mathbb{F}_2^k} (-1)^{|\mathbf{y}|-1} (\chi_{\mathbf{y}}(\mathbf{x}') + \chi_{\mathbf{y}}(\mathbf{x}'') - \chi_{\mathbf{y}}(\mathbf{x}''')) \\ &= \sum_{\substack{\mathbf{y} \in \mathbb{F}_2^k, \\ y_k=0}} (-1)^{|\mathbf{y}|-1} \chi_{\mathbf{y}}(\mathbf{x}') + \sum_{\substack{\mathbf{y} \in \mathbb{F}_2^k, \\ y_k=1}} (-1)^{|\mathbf{y}|-1} (\chi_{\mathbf{y}}(\mathbf{x}') + \chi_{\mathbf{y}}(\mathbf{x}'') - \chi_{\mathbf{y}}(\mathbf{x}''')) \\ &= \sum_{\substack{\mathbf{y} \in \mathbb{F}_2^k, \\ y_k=0}} (-1)^{|\mathbf{y}|-1} (\chi_{\mathbf{y}}(\mathbf{x}') - \chi_{\mathbf{y}}(\mathbf{x}') \oplus x_k - \chi_{\mathbf{y}}(\mathbf{x}') \oplus x_{k+1} + \chi_{\mathbf{y}}(\mathbf{x}') \oplus x_k \oplus x_{k+1}) \\ &= \sum_{\mathbf{y} \in \mathbb{F}_2^{k+1}} (-1)^{|\mathbf{y}|-1} \chi_{\mathbf{y}}(\mathbf{x}) \end{aligned}$$

□

Corollary B.2. *For any $\mathbf{x}, \mathbf{y} \in \mathbb{F}_2^n$,*

$$2^{|\mathbf{y}|-1} \mathbf{x}^{\mathbf{y}} = \sum_{\mathbf{y}' \subseteq \mathbf{y}} (-1)^{|\mathbf{y}'|-1} \chi_{\mathbf{y}'}(\mathbf{x})$$