

Efficient quantum programming using EASE gates on a trapped-ion quantum computer

Nikodem Grzesiak¹, Andrii Maksymov¹, Pradeep Niroula^{2,3}, and Yunseong Nam^{1,4}

¹IonQ, College Park, MD 20740, USA

²Joint Center for Quantum Information and Computer Science, NIST/University of Maryland, College Park, MD 20742, USA

³Joint Quantum Institute, University of Maryland, College Park, MD 20742, USA

⁴Department of Physics, University of Maryland, College Park, MD 20742, USA

Parallel operations in conventional computing have proven to be an essential tool for efficient and practical computation, and the story is not different for quantum computing. Indeed, there exists a large body of works that study advantages of parallel implementations of quantum gates for efficient quantum circuit implementations. Here, we focus on the recently invented efficient, arbitrary, simultaneously entangling (EASE) gates, available on a trapped-ion quantum computer. Leveraging its flexibility in selecting arbitrary pairs of qubits to be coupled with any degrees of entanglement, all in parallel, we show an n -qubit Clifford circuit can be implemented using $6 \log(n)$ EASE gates, an n -qubit multiply-controlled NOT gate can be implemented using $3n/2$ EASE gates, and an n -qubit permutation can be implemented using six EASE gates. We discuss their implications to near-term quantum chemistry simulations and the state of the art pattern matching algorithm. Given Clifford + multiply-controlled NOT gates form a universal gate set for quantum computing, our results imply efficient quantum computation by EASE gates, in general.

1 Introduction

A gate-based, universal quantum computer is increasingly becoming a commodity between researchers in a variety of fields, and in all instances of use cases, the process of using a quantum computer inevitably includes the key step of compiling a program to an executable targeted to backend quantum hardware. With different architectures available today, leveraging unique capabilities that each different architecture offers is a step to be considered seriously, should one wish to harness the most out of a quantum computer. A parallel may be drawn to conventional computing, where, for example, single instruction multiple data (SIMD) techniques have been used to perform efficient computation [1].

In this paper, we consider a trapped-ion quantum computer that offers a powerful SIMD-like technique called efficient, arbitrary, simultaneously entangling (EASE) gate, as its native operation. As detailed in Sec. 2 (see also [2]), an EASE gate can entangle

Nikodem Grzesiak: nikodem.grzesiak@gmail.com

Andrii Maksymov: andrii.maksymov@gmail.com

Pradeep Niroula: pniroula@umd.edu

Yunseong Nam: ynam@umd.edu

arbitrarily selected pairs of qubits in one step. When the entanglement couplings for all pairs are identical, the EASE gate becomes the so-called global Mølmer-Sørensen (GMS) gate, investigated in [3–5] to demonstrate more efficient compilation of quantum programs when compared to a serial approach.

Specific to this work, we focus on three important classes of circuits, widely used in many quantum programs: Clifford circuits, multiply-controlled NOT ($\text{NOT} := \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$) operation, and the permutation operator. For the first, we consider the number of EASE gates sufficient to implement an n -qubit arbitrary Clifford circuit. Specifically, in Sec. 3, we improve the best-known bound of $6n$, reported in [5], to $6 \log(n)$, where we dropped additive constant dependence. For the second, we consider an n -controlled NOT gate and show, in Sec. 4, the number of EASE gates required can be reduced to $3n/2$. This may be compared to the state-of-the-art $2n$ [4] for the GMS-based approach. Further, we show the number of EASE gates can be as small as two for an arbitrarily large n , should the ancillae be inexpensive. Lastly, in Sec. 5 we show a permutation of arbitrary number of qubits can be performed using six EASE gates only.

2 EASE gates

We start by defining an EASE gate, i.e.,

$$\text{EASE}(\vec{\phi}, \vec{\theta}) := \prod_{j>k} \exp\left(-i\sigma_{\phi_j}^{(j)}\sigma_{\phi_k}^{(k)}\theta_{jk}/2\right), \quad (1)$$

where

$$\sigma_{\phi_j}^{(j)} = \cos(\phi_j)\sigma_x^{(j)} + \sin(\phi_j)\sigma_y^{(j)} \quad (2)$$

is a Pauli operator, defined over a vector that points to the equator of a Bloch sphere with azimuthal angle ϕ_j , acting on qubit j and free parameters θ_{jk} are the entanglement coupling between qubits j and k . Shown in [2] was that, even though the number of θ_{jk} parameters increases quadratically in the number of qubits n , the complexity of the control signal design scales at most linearly in the number of qubits n , bounded from above by $3n - 1$. Note a single two-qubit gate requires similar complexity, i.e., $2n + 1$ [6]. In other words, an EASE gate’s complexity is at most a small (< 1.5) multiplicative constant factor more than that of the conventional two-qubit gate.

It thus stands to reason that the use of EASE gates, whenever possible, should be explored.¹ The main differences between the EASE gates and the GMS gates are that (i) one can choose arbitrary pairs to be entangled in an EASE gate whereas one must entangle all pairs in the GMS gate² and (ii) the entanglement couplings can be chosen flexibly for each selected pair in an EASE gate whereas in the GMS gate they cannot be. In other words, a GMS gate is a special EASE gate with θ_{jk} for all j and k being identical to one another. We note that, in this paper, we judiciously use both (i) and (ii) to improve the quantum circuit efficiency.

Due to their versatile use throughout this paper, we briefly discuss the complexity of implementing a fan-in or a fan-out controlled-NOT (CNOT) operation in the following. It was reported in [3] that two GMS gates are needed to implement either the fan-in or

¹We note that EASE gates could, e.g., be of lower fidelity than the conventional two-qubit gate in today’s trapped-ion quantum computers, as observed in [2]. Therefore, at least for the near term, a care needs to be taken when determining if the use of EASE gates would indeed be beneficial.

²We abuse the notion GMS, including the case where only a subset of qubits is considered.

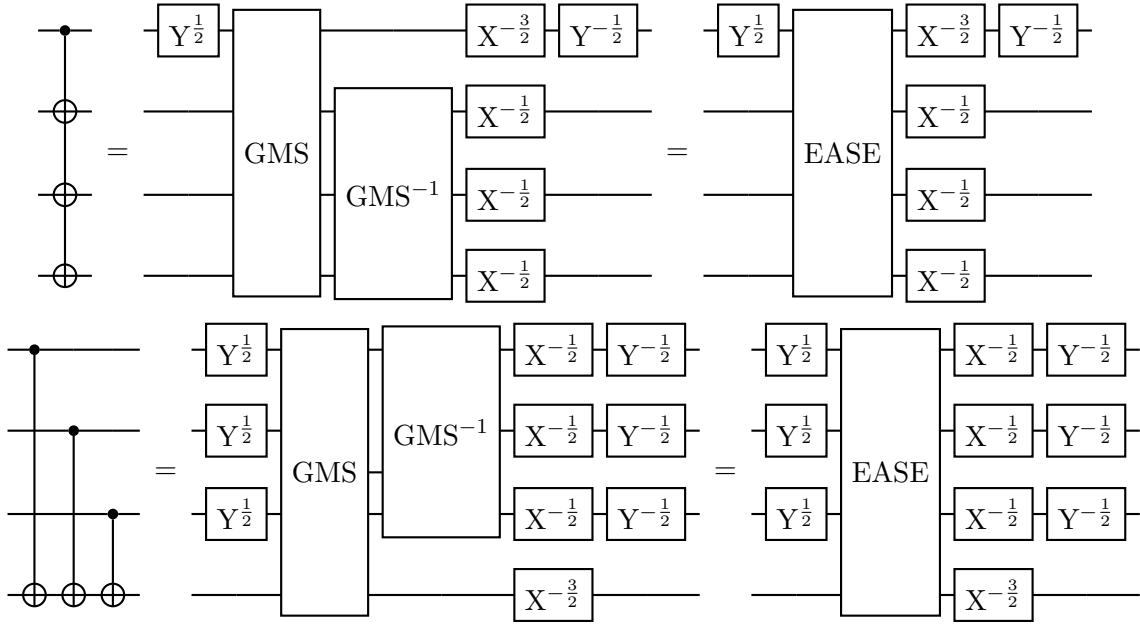


Figure 1: Fan-out and fan-in operation implementation using GMS vs. EASE gates for a four-qubit instance. Both GMS and EASE gates used here assume the entangling angle of $\pi/2$ between pairs of qubits. X^c and Y^c denote the single-qubit operators $\exp(-i\sigma_x c\pi/2)$ and $\exp(-i\sigma_y c\pi/2)$, respectively. The EASE gate used for the fan-out simultaneously entangles only qubits (0, 1), (0, 2), and (0, 3), where qubits are labeled from zero to three from the top. The EASE gate used for the fan-in simultaneously entangles only qubits (0, 3), (1, 3), and (2, 3).

the fan-out operation over n qubits. Note a serial approach would require $n - 1$ $XX(\pi/2)$ gates each, where an $XX(\theta)$ gate is defined as $\exp(-i\sigma_x \sigma_x \theta/2)$. Leveraging the ability to target arbitrary pairs of qubits, only a single EASE gate is required for each. See Fig. 1 for details. Briefly, an EASE gate can implement any combinations of XX gates in parallel, colliding or not at a qubit. All of the XX gates used to implement either the fan-in or the fan-out operation would thus be implemented using a single EASE gate. Parallel fan-ins to different targets with potentially colliding controls or parallel fan-outs from different controls to potentially colliding targets would also cost only a single EASE gate.

3 Clifford circuits

We report in this section an EASE-based method to synthesize an arbitrary Clifford circuit, using a normal form of H-S-CZ-CNOT-H-CZ-S-H [7]. Here, H and S denote single-qubit gate layers of Hadamard $H := \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ and phase gates $S := \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}$, respectively, which are trivially parallelizable. Thus, we focus in the following an efficient method to implement CZ – which stands for controlled-Z, where $Z := \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$, – and CNOT layers.

A CZ layer is straightforward to implement using a single EASE gate. This is so since, as shown in Fig. 2, a CZ gate can be decomposed into a $ZZ := \exp(-i\sigma_z \sigma_z \theta/2)$ gate and two S^{-1} gates, where ZZ and S^{-1} gates commute with one another. This way, all of the ZZ gates can be layered to form a single block of ZZ gates and the aggregated S^{-1} gates, forming a single-qubit gate layer, can be implemented straightforwardly in parallel. The single block of ZZ gates is nothing but a Hadamard-layer conjugated XX gates, and our EASE gates can implement multiple XX gates over arbitrary pairs in a single step.

We next turn to an efficient implementation of a CNOT layer. Discussed in [3] was the




Figure 2: Decomposition of a CZ gate using ZZ and S^{-1} gates, and then into XX, H, and S^{-1} gates. XX and ZZ gates here assume an entanglement angle $\theta = \pi/2$.




Figure 3: Illustration of the implementation strategy of an upper triangular transformation. As detailed in the main text, we first consider the 2×2 regions along the diagonal. For each subsequent iteration, we multiply a factor of two to the sidelengths, exponentially expanding the areas of interest per region.

implementation detail of a CNOT layer, where the layer was LU-decomposed first, then each of the two resulting “triangular” transformations of size $n \times n$ was implemented using $2n$ GMS gates, where two GMS gates were spent for each of the n fan-out CNOTs. This results in $4n$ GMS gates required for a CNOT layer. A naive replacement of two GMS gates used in each fanout with one EASE, possible due to the aforementioned flexibility of the EASE to apply XX gates only between the qubit pairs that comprise of the shared control and all of the targets, would result in $2n$ EASE gates per CNOT layer.

A more efficient approach is to use $n/2$ ancilla qubits, where the ancilla qubits can be used to store XORs of input Boolean variables. To start, note an upper triangular transformation takes the input Boolean variables, say b_i , $i = 0, 1, \dots, n-1$, for each qubit i and outputs $o_i = \sum_{j=i}^{n-1} s_j^{(i)} b_j$, where $s_j^{(i)} \in \{0, 1\}$ is the j th element of the i th row of the linear, triangular transformation matrix. To facilitate the foregoing discussion, we assume n is a power of two, although this is not necessary.

We zoom in on the smallest triangles, regions occupying a 2×2 area each, along the diagonal as shown in Fig. 3, of which there are $n/2$. The upper-right corner of each triangle tells us if the upper element needs to be transformed into the XOR of the upper element and the lower element of the triangle. We achieve this transformation by (a) introducing $n/2$ clean ancilla qubits initialized to $|0\rangle$, (b) computing, in this case copying, the lower element onto the ancilla, should the upper triangular element be one, (c) applying CNOTs, controlled on each ancilla, targeting each upper elements, and (d) uncomputing (b), returning all ancillae to $|0\rangle$.

We next zoom out to focus on $n/4$ triangles along the diagonal, each occupying a 4×4 footprint each (see Fig. 3, middle panel). The idea is similar to that of the 2×2 case considered above, i.e., (a)-(d). This time, each 4×4 includes two 2×2 triangles (shaded region) that we already took care of. Thus, all we need to consider at this point is the transformations implied by the upper-right 2×2 square (light region) of the 4×4 triangle.

Concretely, (a) given two clean ancillae, (b) we first compute, on the ancillae, the XOR values implied by the upper-right 2×2 square using fan-ins from the lower-right 2×2 triangle, of the 4×4 triangle. Since the qubits that correspond to the lower two elements contain XOR patterns of the elements, one can always compute the rows of the square by fan ins. Once the ancillae encode the square, (c) we impart the computed XORs onto the upper two elements by a layer of parallel CNOTs as before. Lastly, (d) we uncompute the ancillae based on the unchanged, lower two elements of each 4×4 triangle.

We iterate the above process by doubling the sidelength of the triangles at each iteration. The number of ancilla qubits required remains constant at $n/2$ and the entire process is over in $\log_2(n)$ steps. Each step requires three EASE gates, one each for steps (b), (c), and (d). Since we have two triangles in the LU decomposition of a CNOT transformation layer in the Clifford normal form, we use, at most, a total of $6 \log_2(n)$ EASE gates for the CNOT stage. Together with the EASE counts required for the two CZ stages, i.e., two, we use for an arbitrary Clifford transformation over n qubits $6 \log_2(n) + 2$ EASE gates, which may be compared to the state-of-the-art $6n$ GMS gates reported in [5].

We note in passing that this construction of a CNOT stage using EASE gates can significantly benefit the near-term application of variational quantum eigensolver, simulating fermionic systems. Shown in [8] was the use of generalized transformation that maps fermion basis to qubit basis to optimize the quantum circuit that results in ground-state energy estimates of a given chemical system. The optimization strategy relies on the ability to induce the transformation efficiently, i.e., the reduction in the circuit complexity due to a good choice of transformation needs to outweigh the resource cost required for implementing the transformation itself in order to gain any advantages. Note the transformation considered therein is precisely the triangular transformation considered in this section. With the EASE approach, the transformation cost is kept at minimum.

4 Multiply-controlled NOT gates

Shown in [3] was the method to construct multiply-controlled, C^{n-1} NOT gates (or Toffoli- n gates, as used in [3]) using GMS gates, where for a $(n-1)$ -control gate, $3n$ GMS gates were used. Later, the state of the art was superseded with the results in [4], requiring only $2n$ GMS gates. In this section, we reduce the number of EASE gates required to $3n/2$. In the extreme, rather unrealistic case of nearly-free ancilla qubits and pulse design resources, we can further reduce the EASE counts to two.

We start our discussion by reiterating the point made in [3], i.e., a C^{n-1} Z gate, which is equivalent to C^{n-1} NOT gate up to a conjugation by Hadamard gates on the target qubit, performs

$$C^{n-1}Z|b_0, b_1, \dots, b_{n-1}\rangle \mapsto (w_{2^n})^{2^{n-1} \prod_{j=0}^{n-1} b_j} |b_0, b_1, \dots, b_{n-1}\rangle, \quad (3)$$

where $w_{2^n} := \exp(i\pi/2^{n-1})$ and we assigned qubit indices j to each Boolean input b_j . Using $2xy = x + y - (x \oplus y)$, it was shown in [3] that the exponent $2^{n-1} \prod_{j=0}^{n-1} b_j$ can be expanded according to

$$2^{n-1} \prod_{j=0}^{n-1} b_j = \sum_{l=1}^n (-1)^{l-1} T_l, \quad (4)$$

where

$$T_l = \sum_{k=1}^{nC_l} \bigoplus_{m=1}^l b_{c_l(k,m)} \quad (5)$$




Figure 4: Circuit diagram relevant for an EASE-based implementation of a C^4 NOT gate. The circuit shown implements the phase inductions detailed in the main text. A C^4 NOT gate is implemented by conjugating the circuit by circuits shown in Fig. 6 in Appendix A, then further conjugating the resulting circuit by Hadamard gates on the target qubit (unspecified). Shown are the single-qubit Z^c gates, which appear in solid boxes, and the ZZ gates, which appear as solid lines, each connecting two qubits marked by solid circles. Dashed boxes are drawn to delineate different pattern lengths l . The entanglement angles θ of ZZ gates for the different l values appear at the bottom of the dashed boxes. Input Boolean variables appear to the left of the circuit.

is the sum of all distinct length- l XOR patterns of input Boolean values, $c_l(k, m)$ is the m 'th qubit index that appears in the k th length- l pattern, and nC_l is n choose l . Note a ZZ(θ) gate on inputs a and b induces, up to a global phase,

$$|ab\rangle \mapsto e^{i\theta(a\oplus b)}|ab\rangle. \quad (6)$$

Given that an EASE gate, conjugated by a Hadamard layer, can implement as many ZZ gates in parallel as one desires with arbitrary angles for each ZZ – in this case $\pm\pi/2^{n-1}$ – the problem of inducing a C^{n-1} NOT gate reduces to inducing Boolean values such that XORs of the pairs of the values span the entire space of T_l for $l = 1, 2, \dots, n$.

We note that, whenever possible, we take advantage of the fact that an application of an Z^c , which is equivalent to $\exp(-ic\pi\sigma_z/2)$ up to a global phase, induces

$$|a\rangle \mapsto e^{i\theta a}|a\rangle, \quad (7)$$

where $\theta = c\pi$. For instance, we take care of all $l = 1$ terms by simply applying appropriate Z^c gates to each qubit. We use this on some of the computed XOR patterns, to be discussed below.

A three-GMS construction of C^2 NOT and C^3 NOT are already available in [3]. We therefore focus on C^4 NOT and C^5 NOT and show that each of them can be constructed with three EASE gates, consuming four and seven ancillae each. To see how the construction works, we assume we start with the input Boolean variables b_j 's. For C^4 NOT, we induce (or compute – used interchangeably for the Boolean variables) $b_2 \oplus b_3, b_2 \oplus b_4, b_3 \oplus b_4$, and $b_0 \oplus b_1 \oplus b_2 \oplus b_3 \oplus b_4$ on to each of the four ancillae (see Fig. 6 in Appendix A). For C^5 NOT, we induce $b_0 \oplus b_1 \oplus b_2 \oplus b_3 \oplus b_4 \oplus b_5, b_0 \oplus b_1 \oplus b_2 \oplus b_5, b_0 \oplus b_1 \oplus b_3 \oplus b_4, b_0 \oplus b_2 \oplus b_3 \oplus b_4, b_1 \oplus b_2 \oplus b_3 \oplus b_4, b_3 \oplus b_5$, and $b_4 \oplus b_5$ on to each of the seven ancillae (see Fig. 7 in Appendix A). It is then possible to show by a simple exhaustive search that an appropriately selected set of pairs of qubits can induce all of $T_l, l = 2, 3, 4$ for C^4 NOT (Fig. 4). For C^5 NOT, the set can cover

all T_l , $l = 2, 3, 4, 5$, except for $b_0 \oplus b_1 \oplus b_2 \oplus b_5$, $b_0 \oplus b_1 \oplus b_3 \oplus b_4$, $b_0 \oplus b_2 \oplus b_3 \oplus b_4$, $b_1 \oplus b_2 \oplus b_3 \oplus b_4$, which are the XOR patterns already encoded in the ancillae (Fig. 5). Therefore, we take care of these four patterns by simply applying appropriate Z^c gates. Similarly, we take care of T_5 for C^4 NOT and T_6 for C^5 NOT using Z^c gates applied to the ancillae that already contain the corresponding XOR patterns.

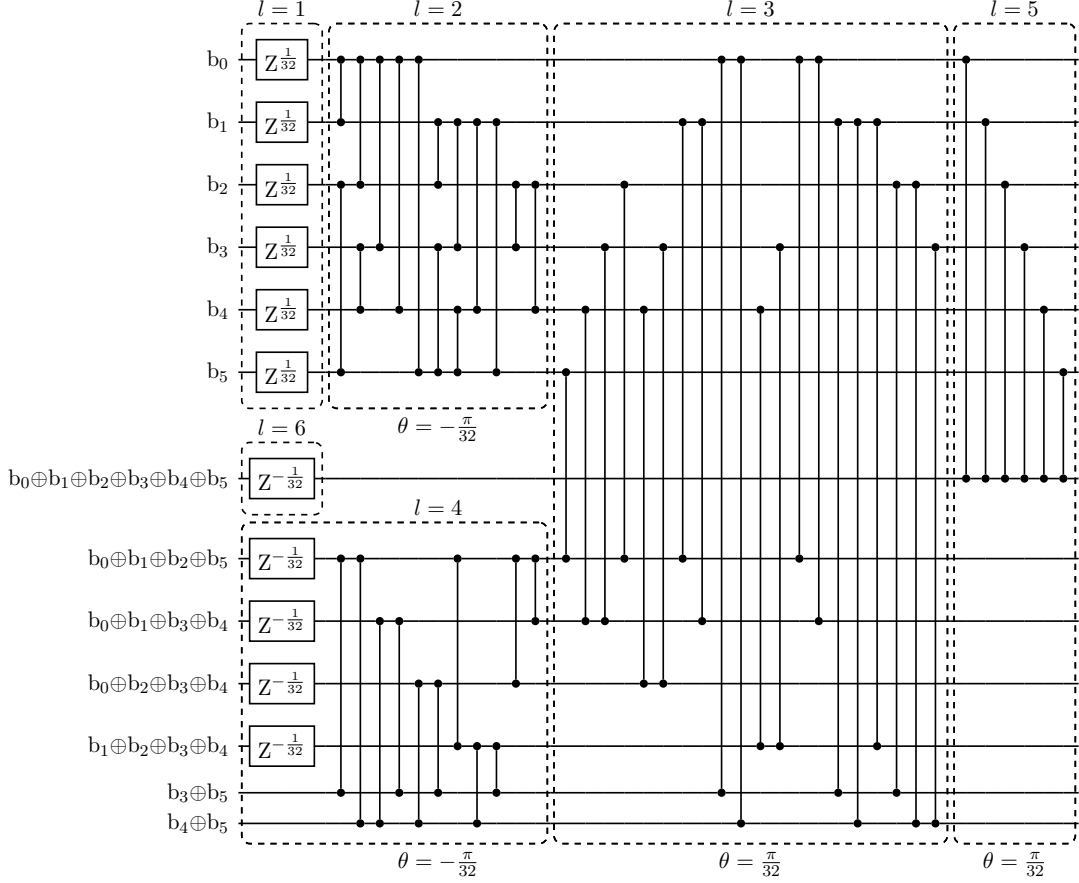


Figure 5: Circuit diagram relevant for an EASE-based implementation of a C^5 NOT gate. The circuit shown implements the phase inductions detailed in the main text. A C^5 NOT gate is implemented by conjugating the circuit by circuits shown in Fig. 7 in Appendix A, then further conjugating the resulting circuit by Hadamard gates on the target qubit (unspecified). The description of the figure is the same as in Fig. 4. For the convenience of the readers we import the caption verbatim. The circuit implements the phase inductions detailed in the main text. Shown are the single-qubit Z^c gates, which appear in solid boxes, and the ZZ gates, which appear as solid lines, each connecting two qubits marked by solid circles. Dashed boxes are drawn to delineate different pattern lengths l . The entanglement angles θ of ZZ gates for the different l values appear at the bottom of the dashed boxes. Input Boolean variables appear to the left of the circuit.

Since the induction of the appropriate XOR patterns onto ancillae cost a single EASE gate, as they are simply a parallel fan-in operations onto each ancilla qubits, computing and uncomputing the patterns cost two EASE gates. Combined with the parallel ZZ operation described above, which costs a single EASE gate, we arrive at a three-EASE construction of C^4 NOT and C^5 NOT.

By use of our C^4 NOT and C^5 NOT gates, one can employ a similar approach explored in [3] to extend the number of controls to an arbitrary number. Note a C^{n-1} NOT essentially computes the AND of the $n - 1$ controls. This can thus be implemented by the use of multiple C^4 NOT and C^5 NOT gates, e.g., by computing and aggregating the AND values

of the subsets of size four or five of the controls. Following the method in [3], we can see that, by use of $C^5\text{NOT}$ gates, roughly four controls can be taken care of each time it is used. Provided that a nested $C^5\text{NOT}$ gates need to be uncomputed, we find $n/2$ $C^5\text{NOT}$ gates are needed for a $C^{n-1}\text{NOT}$. Since in our construction each $C^5\text{NOT}$ requires three EASE gates, overall we require $3n/2$ EASE gates per $C^{n-1}\text{NOT}$, with the ancilla count being $n/4$ up to a small additive constant, no more than seven.

In the extreme (unrealistic) scenario where ancilla qubits are inexpensive and the pulse design with exponential complexity is not an issue, a $C^{n-1}\text{NOT}$ gate, regardless of n , can straightforwardly be implemented using only two EASE gates. To see this, we introduce $2^n - n - 1$ ancilla qubits. For each ancilla qubit we compute $\bigoplus_{m=1}^l b_{c_l(k,m)}$, for all $l > 1$ and k , of which there are $2^n - n - 1$. We can then induce appropriate phase to each XOR pattern by a layer of Z^c gates. We uncompute and free the ancilla qubits. The number of EASE gates required is two, since the compute-uncompute operations require parallel fan-in operations, all controls being the original input, each target being ancilla.

We note in passing that the two-EASE based multiply-controlled NOT gate, like its three-EASE based counterpart, can also be used to extend the number of controls to an arbitrary number. Consider a $C^{b-1}\text{NOT}$ gate, requiring two EASE gates and $2^b - b - 1$ ancilla qubits. A $C^{n-1}\text{NOT}$ gate, built with a plurality of the two-EASE based $C^{b-1}\text{NOT}$ gates, $n \gg b$, would then require, approximately, $4n/(b-2)$ EASE gates and $n/(b-2) + 2^b - b - 1$ ancilla qubits. For $b = 6$, the case we considered earlier, they amount to n EASE gates and $n/4 + 57$ ancilla qubits. The latter number indeed motivates our three-EASE based construction, keeping the ancilla counts modest.

5 Qubit Permutations

Another quantum subroutine where EASE gates offer an improvement over the known state-of-start is the implementation of permutations over qubit registers. It is known, from [9], that any permutation over n qubits can be implemented as four layers of disjoint transpositions with n ancillae and, alternatively, as six layers of disjoint transpositions with no ancillae. Since each transposition, or a SWAP gate, can be implemented as three CNOT gates, each layer of disjoint transpositions can be implemented as three layers of EASE gates. Therefore, any permutation over n qubits can be implemented with $O(1)$ EASE gates.

We next consider a controlled permutation operation. Note that, since the $O(1)$ transpositions used to apply a permutation are disjoint, the $O(1)$ controlled transpositions that comprise a controlled permutation operation are disjoint as well, up to the shared control. To achieve the EASE-based implementation, note that a controlled-transposition gate can be implemented using seven CNOT gates [10]. The CNOT gates, forming seven stages, then comprise a single layer of controlled-transposition gates, together with single qubit gates. The CNOT gates in each CNOT stage are either disjoint (conventional parallel CNOTs) or they collide at the shared control (a fan-out operation). Both of these can be implemented in parallel using a single EASE gate. This implies that a controlled permutation can be implemented using $O(1)$ EASE gates.

Using this technique, we can reduce the complexity of algorithms that use controlled permutations. As an example, we consider the circuit depth of string-matching algorithm [10], which matches a pattern of length M in a text of length N . Our EASE-based approach can reduce the depth from the reported $O\left(\sqrt{N}((\log N)^2 + \log M)\right)$ to $O\left(\sqrt{N}(\log N + \log M)\right)$. The number of ancilla qubits is reduced by $\log N$ as well.

6 Outlook

Harking back to the power and impact of parallel operations in conventional computing, we look forward to the revolution that EASE gates are to bring in quantum computing. Experimental demonstrations have already been successful [2], and the exciting challenges of finding more use cases of EASE gates continue. In this paper, we took a first step towards this goal, leveraging in particular the random access over any pairs of qubits. A future endeavor could include more sophisticated uses of real degrees of freedom in both the amount of entanglement induced between the pairs and the axis of rotation for each Pauli operator acting on each participating qubit in the EASE gates. We hope our work serves as a stepping stone towards efficient quantum computing, adding to the old adage of “circuit depth matters.”

Data availability

All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Information. Additional data related to this paper may be requested from the authors. Correspondence and requests for materials should be addressed to Y.N.(ynam@umd.edu).

Acknowledgements

We thank Qingfeng Wang at the University of Maryland, Christopher Monroe at Duke University, and Igor Markov at the University of Michigan for helpful discussions.

Author contribution

N.G., A.M., and P.N. contributed to devising the methods to optimize various quantum circuit constructions using EASE gates under Y.N.’s supervision. All authors contributed to preparing the manuscript. The authors declare that they have no competing interests.

References

- [1] C. J. Hughes, *Synthesis Lectures on Computer Architecture* **10**, 1 (2015).
- [2] N. Grzesiak, R. Blümel, K. Wright, K. M. Beck, N. C. Pienti, M. Li, V. Chaplin, J. M. Amini, S. Debnath, J.-S. Chen, *et al.*, *Nature communications* **11**, 1 (2020).
- [3] D. Maslov and Y. Nam, *New Journal of Physics* **20**, 033018 (2018).
- [4] K. Groenland, F. Witteveen, K. Schoutens, and R. Gerritsma, *New Journal of Physics* **22**, 063006 (2020).
- [5] J. van de Wetering, *New Journal of Physics* **23**, 043015 (2021).
- [6] S.-L. Zhu, C. Monroe, and L.-M. Duan, *Europhysics Letters (EPL)* **73**, 485–491 (2006).
- [7] R. Duncan, A. Kissinger, S. Perdrix, and J. Van De Wetering, *Quantum* **4**, 279 (2020).
- [8] Q. Wang, M. Li, C. Monroe, and Y. Nam, *Quantum* **5**, 509 (2021).
- [9] C. Moore and M. Nilsson, *SIAM Journal on Computing* **31**, 799 (2001).
- [10] P. Niroula and Y. Nam, *npj Quantum Information* **7**, 1 (2021).

A CNOT networks for multiply-controlled NOT gates

In this section, we show CNOT networks used to induce the input Boolean variables we use for our EASE-based implementation of multiply-controlled NOT gates. Figure 6 shows the network for a C^4 NOT gate. Figure 7 shows the network for a C^5 NOT gate. The two networks are used before and after the phase induction steps discussed in detail in the main text to implement multiply-controlled Z gates of respective sizes, while freeing up the ancillae. The resulting circuits are conjugated by Hadamard gates on the target qubit of choice in order to induce multiply-controlled NOT gates.




Figure 6: Circuit diagram for computing (uncomputing) ancilla qubit states relevant for our C^4 NOT construction.




Figure 7: Circuit diagram for computing (uncomputing) ancilla qubit states relevant for our C^5 NOT construction.