



New Hardness Results for the Permanent Using Linear Optics

Daniel Grier*

Luke Schaeffer†

Abstract

In 2011, Aaronson gave a striking proof, based on quantum linear optics, showing that the problem of computing the permanent of a matrix is $\#P$ -hard. Aaronson's proof led naturally to hardness of approximation results for the permanent, and it was arguably simpler than Valiant's seminal proof of the same fact in 1979. Nevertheless, it did not prove that computing the permanent was $\#P$ -hard for any class of matrices which was not previously known. In this paper, we present a collection of *new* results about matrix permanents that are derived primarily via these linear optical techniques.

First, we show that the problem of computing the permanent of a real orthogonal matrix is $\#P$ -hard. Much like Aaronson's original proof, this will show that even a multiplicative approximation remains $\#P$ -hard to compute. The hardness result even translates to permanents over finite fields, where the problem of computing the permanent of an orthogonal matrix is Mod_pP -hard in the finite field \mathbb{F}_{p^4} for all primes $p \neq 2, 3$. Interestingly, this characterization is tight: in fields of characteristic 2, the permanent coincides with the determinant; in fields of characteristic 3, one can efficiently compute the permanent of an orthogonal matrix by a nontrivial result of Kogan.

Finally, we use more elementary arguments to prove $\#P$ -hardness for the permanent of a positive semidefinite matrix, which shows that certain probabilities of boson sampling experiments with thermal states are hard to compute exactly despite the fact that they can be efficiently sampled by a classical computer.

1 Introduction

The permanent of a matrix has been a central fixture in computer science ever since Valiant showed how it could efficiently encode the number of satisfying solutions to classic NP-complete constraint satisfaction problems [31]. His theory led to the formalization of many of the counting classes in complexity theory, including $\#P$, the complexity class containing those functions which count the number of accepting paths of a polynomial-time non-deterministic Turing machine. Indeed, the power of these counting classes was later demonstrated by Toda's celebrated theorem, which proved that every language in the polynomial hierarchy could be computed in polynomial time with only a single call to a $\#P$ oracle [25].

Let us recall the definition of the matrix permanent. Suppose $A = \{a_{i,j}\}$ is an $n \times n$ matrix over some field. The *permanent* of A is

$$\text{Per}(A) = \sum_{\sigma \in S_n} \prod_{i=1}^n a_{i,\sigma(i)}$$

*MIT. Email: grierd@mit.edu. Supported by an NSF Graduate Research Fellowship under Grant No. 1122374.

†MIT. Email: lrs@mit.edu.

‡Both authors were supported by the Vannevar Bush Faculty Fellowship from the US Department of Defense. Part of this research was completed while visiting UT Austin.

where S_n is the group of permutations of $\{1, 2, \dots, n\}$. Compare this to the *determinant* of A :

$$\det(A) = \sum_{\sigma \in S_n} \text{sign}(\sigma) \prod_{i=1}^n a_{i, \sigma(i)}.$$

Since we can compute the determinant in polynomial time (in fact, in NC^2 ; see Berkowitz [7]), the apparent difference in complexity between the determinant and permanent comes down to the cancellation of terms in the determinant.

In his original proof [31], Valiant casts the permanent in a combinatorial light, in terms of a directed graph rather than as a polynomial. Imagine the matrix A encodes the adjacency matrix of a weighted graph with vertices labeled $\{1, \dots, n\}$. Each permutation σ on the vertices has a cycle decomposition, which partitions the vertices into a collection of cycles known as a *cycle cover*. The weight of a cycle cover is the product of the edge weights of the cycles (i.e., $\prod_{i=1}^n a_{i, \sigma(i)}$), and then the permanent is the sum of the weights of all cycle covers of the graph. Equipped with this combinatorial interpretation of the permanent, Valiant constructs a graph by linking together different kinds of gadgets in such a way that some cycle covers correspond to solutions to a CNF formula, and the rest of the cycle covers cancel out.

Valiant’s groundbreaking proof, while impressive, is fairly opaque and full of complicated gadgets. A subsequent proof by Ben-Dor and Halevi [6] simplified the construction, while still relying on the cycle cover interpretation of the permanent. In 2009, Rudolph [22] noticed an important connection between quantum circuits and matrix permanents—a version of a correspondence we will use often in this paper. Rudolph cast the cycle cover arguments of Valiant into a more physics friendly language, which culminated in a direct proof that the amplitudes of a certain class of universal quantum circuits were proportional to the permanent. If he had pointed out that one could embed $\#\text{P}$ -hard problems into the amplitudes of a quantum circuit, then this would have constituted a semi-quantum proof that the permanent is $\#\text{P}$ -hard.

Finally, in 2011, Aaronson [2] (independently from Rudolph) gave a completely self-contained and quantum linear optical proof that the permanent is $\#\text{P}$ -hard. Aaronson’s proof has a few advantages. First, much of the difficulty of the arguments based on complicated cycle cover gadgets is offloaded onto central, well-known theorems in linear optics and quantum computation. Second, it leads naturally to hardness of approximation results for the matrix permanent. Nevertheless, Aaronson conceded that his proof does not give any new classes of matrices for which the permanent is hard.

Results. In this paper, we refine Aaronson’s linear optical proof technique and show that it *can* give $\#\text{P}$ -hardness results that were previously unknown for permanents of several classes of matrices. Specifically, we say that the *permanent is $\#\text{P}$ -hard* for a class of matrices if all functions in $\#\text{P}$ can be efficiently computed with single-call access to an oracle which computes permanents of matrices in that class. In general, the permanent is hard for a function class A if, given an oracle \mathcal{O} for the permanent, $A \subseteq \text{FP}^{\mathcal{O}[1]}$.

Our main result is a linear optical proof that the permanent of a real orthogonal matrix is $\#\text{P}$ -hard. This will have the consequence that the permanent of matrices in any of the classical Lie groups (e.g., invertible matrices, unitary matrices, symplectic matrices) is also $\#\text{P}$ -hard. This emphasizes one advantage of the linear optical approach over the combinatorial approach: since the matrix is constructed by multiplying matrices, it is much easier to handle groups of matrices and properties which are closed under matrix multiplication. These properties can be very difficult to handle in the cycle cover model. For instance, the matrices which arise from Valiant’s construction may indeed be invertible, but this seems to be more accidental than intentional, and a proof of their invertibility appears nontrivial.

In fact, our approach also reveals a surprising connection between the hardness of the permanent of orthogonal matrices over finite fields and the characteristic of the field. First notice that in fields of characteristic 2, the permanent is equal to the determinant and is therefore efficiently computable. Furthermore, there exists an elaborate yet polynomial time algorithm of Kogan [18] that computes the (orthogonal) matrix permanent over fields of characteristic 3. We prove a sharp dichotomy theorem: for fields of characteristic 2 or 3 there is an efficient procedure to compute orthogonal matrix permanents, and for *all* other primes there exists a finite field¹ of characteristic p for which the permanent of an orthogonal matrix (over that field) is as hard as counting the number of solutions to a CNF formula mod p .² Furthermore, there even exist infinitely many primes for which computing the permanent of an orthogonal matrix over \mathbb{F}_p (i.e., modulo p) is hard.

Finally, we give a polynomial interpolation argument showing that the permanent of a positive semidefinite matrix is $\#\text{P}$ -hard. This has an interesting consequence due a recent connection between matrix permanents and boson sampling experiments with thermal input states [10, 21]. In particular, the probability of detecting one photon on each of a set of modes, and zero photons on the remaining modes, is proportional to a positive semidefinite matrix dependent on the temperatures of the input thermal states. Our result implies that we cannot in general compute such output probabilities exactly despite the fact that an efficient classical sampling algorithm exists [21].

In conclusion, we show that the linear optical framework can help us prove new classical results in computer science. Indeed, we can even use the formalism to prove structural results about the permanent (e.g., increasing the eigenvalues $\lambda_1, \dots, \lambda_n$ of a positive semidefinite matrix $U \text{diag}\{\lambda_1, \dots, \lambda_n\} U^\dagger$ cannot decrease its permanent) for which we currently have no application, but which might later prove useful (see Appendix E for details).

Proof Outline. The main result concerning the $\#\text{P}$ -hardness of real orthogonal permanents follows from three major steps:

1. Construct a quantum circuit (over qubits) with the following property: *If* you can compute the probability of measuring the all-zeros state after the circuit has been applied to the all-zeros state, then you could calculate some $\#\text{P}$ -hard quantity. We must modify the original construction of Aaronson [2], so that all the gates in the used in this construction are real.
2. Use a modified version of the Knill, Laflamme, Milburn protocol [17] to construct a linear optical circuit which simulates the qubit circuit we constructed in the previous step. In particular, we must ensure that the linear optical circuit starts and ends with one photon in every mode. Notice that this is distinct from Aaronson’s approach [2] because we can no longer immediately use the dual-rail encoding of KLM. We build new postselected encoding and decoding gadgets to circumvent this problem.
3. Use the connection between the transition amplitude of a linear optical circuit and the permanent of its underlying matrix which was first pointed out by Caianiello [9] and later used in Aaronson’s proof. Because we paid special attention to the distribution of photons across the modes of our linear optical network, the permanent of the underlying transition matrix will perfectly capture the success probability of the linear optical circuit. It is then simple to work backwards from this permanent to calculate our original $\#\text{P}$ -hard quantity.

¹We can always assume that this field is \mathbb{F}_{p^4} , although in some cases \mathbb{F}_{p^2} or even \mathbb{F}_p will suffice. See Section 4 for more details.

²Formally, this language is complete for the class Mod_kP . A language is in the class Mod_kP if there exists a polynomial-time non-deterministic Turing machine such that a string x is in the language iff the number of accepting paths of the Turing machine is nonzero mod k . By Toda’s theorem, we have that $\text{PH} \subseteq \text{BPP}^{\text{Mod}_k\text{P}}$. See Appendix A for a more precise exposition of such counting classes.

The paper is organized as follows. Section 2 gives a brief introduction to the linear optical framework and the relevant tools we use in this paper. In Section 3, we use this framework to show that the permanent of a real orthogonal matrix is $\#P$ -hard. A careful analysis in Section 4 (and Appendix C) extends these gadgets to finite fields.³ Finally, in Section 5, we explore other matrix classes, culminating in a proof that the permanent of a real special orthogonal symplectic involution is $\#P$ -hard.

2 Linear Optics Primer

In this section we will introduce the so-called “boson sampling” model of quantum computation, which will make clear the connection between the dynamics of *noninteracting* bosons and the computation of matrix permanents [9, 28]. The most promising practical implementations of this model are based on linear optics and use photons controlled by optical elements such as beamsplitters. We will use the term “linear optics” throughout, although the same experiments could be done with other kinds of bosons.

To understand how such particles can have interesting system dynamics, let us discuss the general model of their evolution. Consider first a single boson. At any point in time, it is in one of finitely many *modes*. As the system evolves, the particle moves from one of, say, m initial modes to a superposition of m final modes according to a *transition matrix of amplitudes*. That is, there is an $m \times m$ unitary transition matrix $U \in \mathbb{C}^{m \times m}$, where U_{ji} is the amplitude of a particle going from mode i to mode j .

The model becomes more complex when we consider a system of multiple particles evolving on the same modes according to the same transition matrix. Let us define states in our space of k bosons in what is called the Fock basis. A *Fock state* for a k -photon, m -mode system is of the form $|s_1, s_2, \dots, s_m\rangle$ where $s_i \geq 0$ is the number of bosons in the i th mode and $\sum_{i=1}^m s_i = k$. Therefore, the Hilbert space which spans the Fock basis states $\Phi_{m,k}$ has dimension $\binom{k+m-1}{k}$. Let φ be the transformation which lifts the unitary U to act on a multi-particle system. On a k -particle system, $\varphi(U)$ is a linear transformation from $\Phi_{m,k}$ to $\Phi_{m,k}$.

Let $|S\rangle = |s_1, s_2, \dots, s_m\rangle$ be the Fock state describing the starting state of the system, and let $|T\rangle = |t_1, t_2, \dots, t_m\rangle$ be the ending state. We have:

$$\langle T | \varphi(U) | S \rangle = \frac{\text{Per}(U_{S,T})}{\sqrt{s_1! \dots s_m! t_1! \dots t_m!}}$$

where $U_{S,T}$ is the matrix obtained by taking s_i copies of the i th row and t_i copies of the column i in U for all $i \in \{1, 2, \dots, m\}$. We will refer to this formula as the φ -*transition formula*. Notice that $s_1 + \dots + s_m$ must equal $t_1 + \dots + t_m$ in order for $U_{S,T}$ to be square. This expresses the physical principle that photons are not created or destroyed in the experiment.

For example, suppose U is the common Hadamard gate and that we wish to apply U to two modes each with a single photon. That is, $U = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ and $|S\rangle = |1, 1\rangle$. Since the number of photons must be conserved, the resulting state of the system is in some linear combination of $|2, 0\rangle$, $|1, 1\rangle$, and $|0, 2\rangle$. We calculate these amplitudes explicitly below:

³As is the case with Aaronson’s proof, our real orthogonal construction also leads naturally to hardness of approximation results, which we discuss in Appendix D.

$ T\rangle$	$ 2, 0\rangle$	$ 1, 1\rangle$	$ 0, 2\rangle$
$U_{S,T}$	$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$	$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -1 \\ 1 & -1 \end{pmatrix}$	$\frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ -1 & -1 \end{pmatrix}$
$\text{Per}(U_{S,T})$	1	0	-1
$\langle T \varphi(U) S\rangle$	$1/\sqrt{2}$	0	$-1/\sqrt{2}$

Therefore, when we apply Hadamard in a linear optical circuit to the state $|1, 1\rangle$ we get the state $\frac{|2,0\rangle - |0,2\rangle}{\sqrt{2}}$. Indeed, we have derived the famous Hong-Ou-Mandel effect [15]. Although the photons are noninteracting, the fact they are identical results in a final state which is clearly highly entangled. Notice that even when U is unitary, $U_{S,T}$ need not be. However, an important (yet simple) observation for the remainder of this paper is the following:

Observation 1. *If U is unitary and $|S\rangle = |T\rangle = |1, \dots, 1\rangle$, then the transition amplitude from S to T is equal to the permanent of U .*

Proof. If $|S\rangle = |T\rangle = |1, \dots, 1\rangle$, then $U_{S,T} = U$ and $\langle T|\varphi(U)|S\rangle = \text{Per}(U)$. \square

Let us quickly discuss the intuition behind the φ -transition formula. Suppose we have k photons where the i th photon is in mode $x_i \in \{1, \dots, m\}$. We get a contribution to the amplitude of the final state $|T\rangle$ for each function $f : \{1, \dots, m\}^k \rightarrow \{1, \dots, m\}^k$ which maps our initial state into a state consistent with $|T\rangle$. That contribution is simply the product of the amplitudes dictated by our underlying transition matrix: $\prod_{i=1}^k U_{f(x_i), x_i}$. It is not hard to see that summing over all such f is equivalent to computing the function $\text{Per}(U_{S,T})$. The normalization factor $\sqrt{s_1! \dots s_m! t_1! \dots t_m!}$ which we see in the φ -transition formula comes from the fact that we have a system of indistinguishable particles. That is, the permutation which, say, permutes two photons in the same mode, is identical to same permutation which does not permute those two photons.⁴

Finally, we note that φ has the two properties that one would expect from a transformation describing a physical system:

Theorem 2 (see, for example, Aaronson and Arkhipov [2, 3]). *The map φ is a group homomorphism. Furthermore, if $U \in \mathbb{C}^{n \times n}$ is unitary, then $\varphi(U)$ is unitary.*

For our purposes, one of the most important results in linear optics is the following theorem:

Theorem 3 (Knill, Laflamme, and Milburn [17]). *Postselected linear optical circuits are universal for quantum computation. Formally, given a quantum circuit Q consisting of CSIGN and single-qubit gates, there exists a linear optical network U constructible in polynomial time such that*

$$\langle 0, 1, 0, 1 \dots 0, 1 | \varphi(U) | 0, 1, 0, 1 \dots 0, 1 \rangle = \frac{1}{4^\Gamma} \langle 0 \dots 0 | Q | 0 \dots 0 \rangle,$$

where Γ is the number of CSIGN gates in Q .

We will refer to the construction of the linear optical network U from Q in Theorem 3 as the *KLM protocol*. Because it will be necessary to modify the KLM protocol for our reduction in Section 3, it will be helpful to give some idea of its proof here. First, each qubit of Q is encoded in two modes of U in the classic dual-rail encoding. That is, the qubit state $|0\rangle$ is encoded by the Fock state $|0, 1\rangle$ and the state $|1\rangle$ is encoded by the Fock state $|1, 0\rangle$. We will sometimes write $(i, 0)$ and $(i, 1)$ as labels for the first and second mode, respectively, of the dual-rail encoding of the i th qubit. That is, $(i, 0) = |1\rangle$ iff the i th qubit is in state $|1\rangle$.

⁴The square root which appears in the φ -transition formula arises due to the fact that we are computing amplitudes and not probabilities.

Now suppose G is a single-qubit gate applied to qubit i in Q . Using the φ -transition formula, it is not hard to see that applying G to modes $(i, 0)$ and $(i, 1)$ in the linear optical circuit implements the correct single-qubit transformation on the i th encoded qubit. Applying a CSIGN gate is trickier. The KLM protocol builds the CSIGN gate from a simpler NS₁ gate, which flips the sign of a single mode if it has 2 photons and does nothing when the mode has 0 or 1 photon. Using two NS₁ gates one can construct a CSIGN gate. See Figure 6 in Appendix F.

Unfortunately, it is not hard to prove that such a one-mode linear optical gate does not exist. Therefore, some additional resource is required. The original KLM protocol uses *adaptive measurements*, that is, the ability to measure in the Fock basis in the middle of a linear optical computation and adjust the remaining sequence of gates if necessary. Intuitively, using adaptive measurements one can apply some transformation and then measure a subset of the modes to “check” if the NS₁ was applied.

For simplicity, however, we will assume we have a stronger resource—namely, *postselection*. That is, for any measurement counting the number of photons in a particular set of modes, if there is an outcome with nonzero probability, then we can assume that the measurement gives that outcome. With postselection, we can implement an NS₁ gate as follows. First, apply a larger three mode gate G to the state $|x, 0, 1\rangle$ where x is the desired input to NS₁. If one were to measure the last two modes after applying G , then with probability 1/4 the resulting state would be $|0, 1\rangle$. Furthermore, if the measurement outcome is $|0, 1\rangle$, then the operation of G on the first mode is NS₁. Therefore, postselecting on the state $|0, 1\rangle$ guarantees that NS₁ will be applied as desired. Putting the above parts together completes the proof Theorem 3.

3 Permanents of Real Orthogonal Matrices

The first class of matrices we consider are the real orthogonal matrices. That is, square matrices $M \in \mathbb{R}^{n \times n}$ with $MM^T = M^T M = I$. These matrices form a group under composition, the *real orthogonal group*, usually denoted $O(n, \mathbb{R})$. This is a subgroup of the special unitary group, $U(n, \mathbb{C})$, which is itself a subgroup of the general linear group $GL(n, \mathbb{C})$. In other words, our hardness results will carry over to unitary matrices and invertible matrices.⁵

Our result follows the outline of Aaronson’s linear optical proof that the permanent is #P-hard [2]. In particular, our result depends on the KLM construction [17], and a subsequent improvement by Knill [16], which will happen to have several important properties for our reduction. Let us briefly summarize Aaronson’s argument. Suppose we are given a classical circuit C , and wish to compute Δ_C , the number of satisfying assignments minus the number of unsatisfying assignments. Clearly, calculating Δ_C is a #P-hard problem. The first thing to notice is that there exists a simple quantum circuit Q such that the amplitude $\langle 0 \cdots 0 | Q | 0 \cdots 0 \rangle$ is proportional to Δ_C . The KLM protocol [17] implies that there exists a postselected linear optical experiment simulating Q . Therefore, there is some unitary U implemented by a linear optical network, and Fock states $|S\rangle, |T\rangle$ such that the amplitude on $|T\rangle$ when the initial state of the system is $|S\rangle$ is proportional to Δ_C . Since this amplitude is *also* proportional to the permanent of the submatrix $U_{S,T}$ by the φ -transition formula, computing $\text{Per}(U_{S,T})$ allows us to calculate Δ_C .

Notice that Aaronson’s result does *not* imply that the permanent of a unitary matrix is #P-hard. Indeed, it does not even prove that the permanent of an invertible matrix is #P-hard because $U_{S,T}$ need not be invertible when U is invertible.

Indeed, in Aaronson’s proof $|S\rangle = |T\rangle = |0, 1, 0, 1, \dots, 0, 1\rangle$, which comes directly from the dual-rail encoding of the original KLM construction. That is, the submatrix $U_{S,T}$ is one in which

⁵See Corollary 17 for a complete list of classical Lie groups for which our result generalizes.

only alternating rows and columns of the original matrix U are used. Imagine instead that $|S\rangle = |T\rangle = |1, \dots, 1\rangle$. In this case, the matrix $U_{S,T} = U$, and the $\#\text{P}$ -hardness of the permanent of a unitary matrix would follow. In fact, this is *exactly* what we will arrange. Furthermore, we will be careful to use *real* matrices exclusively during all gadget constructions. This will imply that the permanent of real orthogonal matrices is $\#\text{P}$ -hard.

In the following subsections, we will focus on the exact details of the reduction and emphasize those points where our construction differs from that of Aaronson.

3.1 Constructing the Quantum Circuit

Let $C : \{0, 1\}^n \rightarrow \{0, 1\}$ be a classical circuit of polynomial size and let

$$\Delta_C := \sum_{x \in \{0,1\}^n} (-1)^{C(x)}.$$

In this section, we prove the following:

Theorem 4. *Given C , there exists a $p(n)$ -qubit quantum circuit Q such that*

$$\langle 0|^{\otimes p(n)} Q |0\rangle^{\otimes p(n)} = \frac{\Delta_C}{2^n}$$

where $p(n)$ is some polynomial in n . Furthermore, Q can be constructed in polynomial time with a polynomial number of real single-qubit gates and CSIGN gates.

Let \mathcal{O}_C be the standard oracle instantiation of C on $n + 1$ qubits. That is, suppose $\mathcal{O}_C|x, b\rangle = |x, b \oplus C(x)\rangle$ for all $x \in \{0, 1\}^n$ and $b \in \{0, 1\}$. Let us assume we have a Toffoli gate for now and we will see how to construct it later. Since C is a classical boolean function of polynomial complexity, C can be implemented [27] with a polynomial number of Toffoli and NOT gates⁶ and a polynomial number of ancillas starting in the $|0\rangle$ state. Let us describe, briefly, one way to construct \mathcal{O}_C .

Suppose we are given the circuit C as a network of polynomially many NAND gates. For each wire, with the exception of the input wires, we create an ancilla initially in state $|0\rangle$ and use the NOT gate to put it in state $|1\rangle$. For each NAND gate (in topological ordering, i.e., such that no gate is applied before its inputs have been computed), we apply a Toffoli gate targeting the ancilla associated with the output wire, and controlled by the qubits associated with its input wires (whether they are the output of an earlier NAND gate, or an actual input). Hence, the target qubit is in state $|1\rangle$ unless both control qubits are in state $|1\rangle$, simulating a NAND gate. When we have applied all the gates, the output of the function will exist in the final register. We apply a $Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ gate to this register, then apply the sequence of gates in reverse to uncompute the function.

We can now describe the entire circuit for Q (ignoring the ancilla qubits), and it is depicted in Figure 1.

Then,

$$\langle 0|^{\otimes p(n)} Q |0\rangle^{\otimes p(n)} = \frac{1}{2^n} \left(\sum_{x \in \{0,1\}^n} \langle x | \langle - | \right) \mathcal{O}_C \left(\sum_{x \in \{0,1\}^n} |x\rangle |-\rangle \right) = \frac{\Delta_C}{2^n}.$$

Recall that the KLM construction allows the conversion of any quantum circuit into a linear optical circuit *provided* it is constructed from CSIGN and single-qubit gates. Since Q is constructed

⁶Because we require that all ancillas start in the $|0\rangle$ state, we also need the NOT gate to create $|1\rangle$ ancillas.

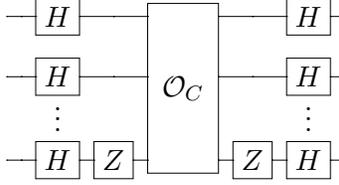


Figure 1: Circuit for Q , where H is the standard Hadamard gate and Z is the Pauli σ_Z matrix.

from Toffoli and single-qubit gates, it suffices to show how the Toffoli gate can be constructed from single-qubit gates and CSIGN gates. Unfortunately, Aaronson’s proof [2] uses a classic construction of the Toffoli gate which uses complex single-qubit gates (see, for example, Nielsen and Chuang [20]), which will later give rise to linear optical circuits with complex matrix representations.⁷ Therefore, we will restrict ourselves to CSIGN and *real* single-qubit gates in our construction of the Toffoli gate.⁸

Lemma 5. *There exists a circuit of CSIGN, Hadamard, and $R_{\pi/4}$ gates which implements a Toffoli gate exactly, where*

$$R_{\pi/4} = \frac{1}{2} \begin{pmatrix} \sqrt{2 + \sqrt{2}} & -\sqrt{2 - \sqrt{2}} \\ \sqrt{2 - \sqrt{2}} & \sqrt{2 + \sqrt{2}} \end{pmatrix}.$$

We prove this lemma in Appendix B. This completes the proof of Theorem 4.

3.2 Postselected Linear Optical Gadgets

We will construct a postselected linear optical circuit O which will simulate the qubit circuit Q on the all zeros input via a modified version of the KLM protocol. Importantly, we will stipulate that our linear optical experiment starts and ends with a single photon in each mode so that in the φ -transition formula we take the permanent of O , not a submatrix of O . Therefore, we will arrive at the following consequence⁹

$$\text{Per}(O) = \langle 1, \dots, 1 | \varphi(O) | 1, \dots, 1 \rangle \propto \langle 0 \cdots 0 | Q | 0 \cdots 0 \rangle \propto \Delta_C.$$

The first step was to convert from a classical circuit to a quantum circuit. Below we formalize the second step: converting from a quantum circuit to a linear optical circuit.

Theorem 6. *Given an n -qubit circuit Q of $\text{poly}(n)$ CSIGN and real single-qubit gates, there exists a linear optical circuit O on $4n + 2\Gamma$ modes such that*

$$\langle 1, \dots, 1 | \varphi(O) | 1, \dots, 1 \rangle = \left(\frac{1}{3} \sqrt{\frac{2}{3}} \right)^\Gamma \left(\frac{-1}{\sqrt{6}} \right)^n \langle 0 |^{\otimes n} Q | 0 \rangle^{\otimes n},$$

where Γ is the number of CSIGN gates in Q . Furthermore, the matrix representing O is a real orthogonal matrix and can be computed in time polynomial in n .

⁷Actually, the proof of Aaronson [2] claims that the final linear optical matrix consists entirely of real-valued entries even though the matrices of the individual single-qubit gates have complex entries. In fact, the matrix *does* have complex entries, but our construction for Toffoli suffices to fix this error.

⁸Although one can argue on abstract grounds that the Toffoli gate and the set of real single-qubit gates suffice to densely generate the orthogonal matrices (i.e., $O(2^n)$ for every $n > 0$) [23], it will turn out to be both simpler and necessary to have an exact decomposition. In particular, we will need an exact construction of the Toffoli gate in Section 4 where we discuss the computation of permanents in finite fields.

⁹To clarify, $|0 \cdots 0\rangle$ is a tensor product of qubits in the state $|0\rangle$ and $|1, \dots, 1\rangle$ is a Fock state with 1 photon in every mode.

We now show how to construct such a linear optical circuit O using the original KLM protocol, subsequent improvements by Knill [16], and a new gadget unique to our problem. First, let us recall our main issue with using the original KLM protocol: to prove that *orthogonal* matrices are #P-hard, we must have that all modes start and end with exactly one photon. There are two instances in which the original KLM protocol requires a mode to be empty at the beginning and end of the computation. First, the NS_1 gate postselects on the Fock state $|0, 1\rangle$, and second, KLM work in a dual-rail encoding. Therefore, half of the modes in the original KLM protocol start and end empty.

To overcome the first obstacle, we appeal to subsequent work of Knill [16], in which the NS_1 gadget construction for CSIGN is replaced by a single 4-mode gadget V , which directly implements CSIGN with two modes postselected in state $|1, 1\rangle$. The matrix for V is the following:

$$V = \frac{1}{3\sqrt{2}} \begin{pmatrix} -\sqrt{2} & -2 & 2 & 2\sqrt{2} \\ 2 & -\sqrt{2} & -2\sqrt{2} & 2 \\ -\sqrt{6+2\sqrt{6}} & \sqrt{6-2\sqrt{6}} & -\sqrt{3+\sqrt{6}} & \sqrt{3-\sqrt{6}} \\ -\sqrt{6-2\sqrt{6}} & -\sqrt{6+2\sqrt{6}} & -\sqrt{3-\sqrt{6}} & -\sqrt{3+\sqrt{6}} \end{pmatrix}$$

from which we can calculate the exact transition amplitudes given below:

$$\begin{aligned} \langle 0, 0, 1, 1 | \varphi(V) | 0, 0, 1, 1 \rangle &= \frac{1}{3} \sqrt{\frac{2}{3}} & \langle 0, 1, 1, 1 | \varphi(V) | 1, 0, 1, 1 \rangle &= 0 \\ \langle 0, 1, 1, 1 | \varphi(V) | 0, 1, 1, 1 \rangle &= \frac{1}{3} \sqrt{\frac{2}{3}} & \langle 1, 0, 1, 1 | \varphi(V) | 0, 1, 1, 1 \rangle &= 0 \\ \langle 1, 0, 1, 1 | \varphi(V) | 1, 0, 1, 1 \rangle &= \frac{1}{3} \sqrt{\frac{2}{3}} & \langle 2, 0, 1, 1 | \varphi(V) | 1, 1, 1, 1 \rangle &= 0 \\ \langle 1, 1, 1, 1 | \varphi(V) | 1, 1, 1, 1 \rangle &= -\frac{1}{3} \sqrt{\frac{2}{3}} & \langle 0, 2, 1, 1 | \varphi(V) | 1, 1, 1, 1 \rangle &= 0 \end{aligned}$$

We now argue that these transition amplitudes suffice to generate a postselected CSIGN. Consider the circuit depicted in Figure 2. First of all, because we postselect on the final two modes ending in the state $|1, 1\rangle$, we only need to consider those transitions for which those two modes end in that state. Secondly, because we use “fresh” ancillary modes for every CSIGN gate, we can always assume that those two modes start in the $|1, 1\rangle$ state. This already vastly reduces the number of cases we must consider. Finally, we wish to know what will happen when the first two modes start in the states $|0, 0\rangle$, $|0, 1\rangle$, $|1, 0\rangle$, and $|1, 1\rangle$. Indeed, our construction will ensure that there is never more than one photon per mode representing one of the dual-rail encoded qubits.

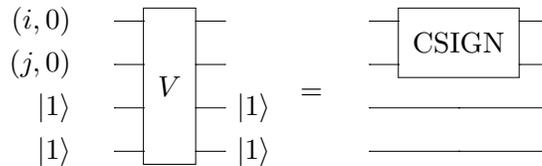


Figure 2: Applying V gadget to generate CSIGN.

For instance, the transition amplitudes of V show that whenever the modes $(i, 0)$ and $(j, 0)$ each start with a photon, there is 0 probability (after postselection) that those photons transition to a state in which one of those modes contains 2 photons and the other contains no photons. Finally, we find that all other amplitudes behave exactly as we would expect for CSIGN. Since each of the acceptable transitions (e.g. from the state $|0, 1\rangle$ to the state $|0, 1\rangle$) has equal magnitude, we only have left to check that V flips the sign of the state whenever the i th and j th qubits are

both in the $|1\rangle$ state. Since this is indeed the case, V implements a postselected CSIGN gate. Importantly, because φ is a homomorphism, we can analyze each such gate separately. Therefore, using the above we can now construct a linear optical circuit O where all of our postselected modes for CSIGN start and end in the $|1\rangle$ state.

We now turn our attention to the dual-rail encoding. Instead of changing the dual-rail encoding of the KLM protocol directly, we will start with one photon in every mode and apply a linear optical gadget to convert to a dual-rail encoding. Of course, the number of photons in the circuit must be conserved, so we will dump these extra photons into n modes separate from the modes of the dual-rail encoding. Specifically, each logical qubit in our scheme is initially represented by four modes in the state $|1, 1, 1, 1\rangle$. We construct a gadget that moves a photon from the first mode to the third mode, postselecting on a single photon in the last mode. That is, under postselection, we get the transition

$$|1, 1, 1, 1\rangle \rightarrow |0, 1, 2, 1\rangle,$$

where the first two modes now represent a $|0\rangle$ qubit in the dual-rail encoding, the third mode stores the extra photon (which we will reclaim later), and the last qubit is necessary for postselection. We call the gadget for this task the *encoding gadget* E , and it is applied to the first, third, and fourth mode of the above state. The matrix for E is

$$E = \frac{1}{\sqrt{6}} \begin{pmatrix} \sqrt{2} & -\sqrt{2} & \sqrt{2} \\ 0 & \sqrt{3} & \sqrt{3} \\ -2 & -1 & 1 \end{pmatrix}$$

from which we get the following transition amplitudes

$$\langle 1, 1, 1 | \varphi(E) | 1, 1, 1 \rangle = 0, \quad \langle 2, 0, 1 | \varphi(E) | 1, 1, 1 \rangle = 0, \quad \langle 0, 2, 1 | \varphi(E) | 1, 1, 1 \rangle = \frac{1}{\sqrt{3}}.$$

After applying the encoding gadget to each logical qubit, we can implement the KLM protocol as previously discussed.¹⁰ Therefore, the relevant amplitude in the computation of Q is now proportional to amplitude of the Fock state which has n groups of modes in the state $|1, 0, 2, 1\rangle$ and 2Γ modes in the state $|1\rangle$. Because we want to return to a state which has one photon in every mode, we must reverse the encoding step.¹¹ For this purpose, we construct a *decoding gadget* D , which will not require any extra postselected modes. We apply the gadget to the second and third modes of the logical qubit such that the two photons in the third mode split with some nonzero probability. The matrix for D is

$$D = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$$

from which the transition condition $\langle 1, 1 | \varphi(D) | 0, 2 \rangle = -1/\sqrt{2}$ follows. Nearly any two-mode linear optical gate would suffice here, but D , the familiar Hadamard gate, optimizes the amplitude above. Notice that if the logical qubit had been in the $|1\rangle$ state, then we would have applied D to a three-photon state. Because of conservativity, the resulting amplitude on the $|1, 1\rangle$ state would be zero. Putting together the encoding gadget, the KLM scheme, and the decoding gadget completes the proof of the theorem.

¹⁰One might wonder why we cannot simply apply the encoding gadget to the *entire* input, thus circumventing the need to use Knill's more complicated V gadget to implement CSIGN. Examining Theorem 3 carefully, we see that all the postselection actually happens at the end of the computation. One might be concerned that once we measured the state $|0, 1\rangle$ to implement NS_1 , those modes would remain in that state. Nevertheless, it *is* possible to compose the gadgets in such a way to allow for postselection on $|0\rangle$ while maintaining that the desired amplitude is still on the $|1, \dots, 1\rangle$ state. We omit such a design since V will turn out to have some nice properties, including its minimal usage of ancillary modes.

¹¹Notice that postselection was required for the encoding gadget, so it does not have a natural inverse.

3.3 Main Result

We are finally ready to prove the main result:

Theorem 7. *The permanent of a real orthogonal matrix is #P-hard. Specifically, given a polynomially sized Boolean circuit C , there exist integers $a, b \in \mathbb{Z}$ and a real orthogonal matrix O computable in polynomial time such that*

$$\text{Per}(O) = 2^a 3^b \Delta_C.$$

Proof. We reduce from the problem of calculating Δ_C for some polynomially sized Boolean circuit C on n bits. By Theorem 4, we first construct the quantum circuit Q from CSIGN and single-qubit gates such that $\langle 0|^{\otimes p(n)} Q |0\rangle^{\otimes p(n)} = \Delta_C / 2^n$. Let Γ be the number of CSIGN gates in Q . We then convert the qubit circuit Q to a linear optical circuit O on $4p(n) + 2\Gamma$ modes using Theorem 6. Notice that we can assume without loss of generality that $p(n)$ and Γ are both even since we can always add an extra qubit to the circuit Q and/or apply an extra CSIGN gate to the $|00\rangle$ state. Combined with the fact that the output amplitudes of linear optical experiments can be described by permanents via the φ -transition formula, we have the following chain of consequences

$$\begin{aligned} \text{Per}(O) &= \langle 1, \dots, 1 | \varphi(O) | 1, \dots, 1 \rangle \\ &= \left(\frac{1}{3} \sqrt{\frac{2}{3}} \right)^\Gamma \left(\frac{-1}{\sqrt{6}} \right)^{p(n)} \langle 0|^{\otimes p(n)} Q |0\rangle^{\otimes p(n)} \\ &= \left(\frac{1}{3} \sqrt{\frac{2}{3}} \right)^\Gamma \left(\frac{-1}{\sqrt{6}} \right)^{p(n)} \left(\frac{1}{2^n} \right) \Delta_C \\ &= 2^a 3^b \Delta_C, \end{aligned}$$

where the last equality comes from the fact that Γ and $p(n)$ are even. □

4 Permanents over Finite Fields

Valiant's foundational work on #P is well-known, but his contemporary work on the relationship between the permanent and the class we now know as Mod_kP is less appreciated. In another 1979 paper [30], Valiant showed that the permanent modulo p is Mod_pP -complete, except when $p = 2$, in which case the permanent coincides with the determinant because $1 \equiv -1 \pmod{2}$.

Theorem 8 (Valiant [30]). *The problem of computing $\text{Per}(M) \pmod{p}$ for a square matrix $M \in \mathbb{F}_p^{n \times n}$ is Mod_pP -complete for any prime $p \neq 2$ (and in NC^2 otherwise).*

As discussed in Appendix A, Mod_pP -hardness provides evidence for the difficulty of computing the permanent, even modulo a prime. In particular, an efficient algorithm for the problem would collapse the polynomial hierarchy.

In the spirit of our result on real orthogonal matrices, we ask whether the permanent is still hard for orthogonal matrices in a finite field. We are not the first to consider the problem; there is the following surprising theorem of Kogan [18] in 1996.

Theorem 9 (Kogan [18]). *Let \mathbb{F} be any field of characteristic 3. There is a polynomial time algorithm to compute the permanent of any orthogonal matrix over \mathbb{F} .*

In other words, for orthogonal matrices, the permanent is easy to compute for fields of characteristic 2 (since it is easy in general), but it is also easy for fields of characteristic 3 (by a much more elaborate argument)! Could it be that the permanent is easy for all finite fields of some other characteristic? No, it turns out. Using the gadgets from Section 3, we prove a converse to Theorem 9.

Theorem 10. *Let $p \neq 2, 3$ be a prime. There exists a finite field of characteristic p , namely \mathbb{F}_{p^4} , such that the permanent of an orthogonal matrix in \mathbb{F}_{p^4} is Mod_pP -hard.*

We prove the theorem by carefully porting Theorem 7 to the finite field setting. Recall that Theorem 7 takes a circuit C and constructs a sequence of gadgets G_1, \dots, G_m such that

$$\text{Per}(G_1 \cdots G_m) = 2^a 3^b \Delta_C, \quad (1)$$

for some $a, b \in \mathbb{Z}$. In general, there is no way to convert such an identity on real numbers into one over finite fields, but all of our gadgets are built out of *algebraic* numbers. In particular, all of the entries are in some algebraic field extension $\mathbb{Q}(\alpha)$ of the rationals, where $\alpha \approx 4.182173283$ is the largest real root of irreducible polynomial

$$f(x) = x^{16} - 40x^{14} + 572x^{12} - 3736x^{10} + 11782x^8 - 17816x^6 + 11324x^4 - 1832x^2 + 1.$$

Each element in $\mathbb{Q}(\alpha)$ can be written as a polynomial (of degree less than 16) in α over the rationals. In Appendix C.1, we give explicit canonical representations for a set of numbers which generate (via addition, subtraction and multiplication, but *not* division) the entries of all our gadgets.

Each entry of a gadget G_i is a polynomial in α with rational coefficients, so observe that we can take a common denominator for the coefficients and write the entry as an integer polynomial divided by some positive integer. By the same token, we can take a common denominator for the entries of a gadget G_i , and write it as $\frac{1}{k_i} \hat{G}_i$ where \hat{G}_i is a matrix over $\mathbb{Z}[\alpha]$, and k_i is a positive integer.

Now we would like to take Equation 1 modulo a prime p . In principle, we can pull k_1, \dots, k_m out of the permanent, multiply through by $Z = (k_1 \cdots k_m)^n 2^{|a|} 3^{|b|}$ to remove all fractions on both sides, and obtain an equation of the form

$$K \text{Per}(\hat{G}_1 \cdots \hat{G}_m) = K' \Delta_C,$$

where K, K' are integers. Then the entire equation is over $\mathbb{Z}[\alpha]$, so if we reduce all the coefficients modulo p , we get an equation over $\mathbb{F}_p[\alpha]$.

We show in Appendix C.1 that for each gadget we use, the denominator k_i may have prime divisors 2, 3, and 23, but no others. Hence, as long as $p \neq 2, 3, 23$ (and in the case $p = 23$, there is an alternative representation we can use, see Appendix C), we can divide through by Z , pull it back inside the permanent as the $\frac{1}{k_i}$ s, and distribute each $\frac{1}{k_i}$ into the corresponding \hat{G}_i . This gives

$$\text{Per}(G_1 \cdots G_m) \equiv 2^a 3^b \Delta_C \pmod{p},$$

the equivalent of Equation 1, but over $\mathbb{F}_p[\alpha]$. In particular, G_1, \dots, G_m are now orthogonal matrices in $\mathbb{F}_p[\alpha]$, and Δ_C has been reduced modulo p .

Note that $\mathbb{F}_p[\alpha] \cong \mathbb{F}_p[x]/(f(x))$ is a ring, not a field. If $f(x)$ were irreducible modulo p then it would be a field, but this will never happen for our f . Consider the following lemma.

Lemma 11. *Let q be a prime power. Suppose \mathbb{F}_q is the subfield of order q contained in the finite field \mathbb{F}_{q^2} . Then every element in \mathbb{F}_q has a square root in \mathbb{F}_{q^2} .*

Proof. Let a be an arbitrary element of \mathbb{F}_q . By definition, a has a square root if the polynomial $f(x) := x^2 - a$ has a root. If f has a root in \mathbb{F}_q then we are done. Otherwise, f is irreducible, but has a root in $\mathbb{F}_q[x]/\langle f(x) \rangle \cong \mathbb{F}_{q^2}$. \square

By Lemma 11, the square roots of 2 and 6 are in \mathbb{F}_{p^2} , and therefore so are $2 + \sqrt{2}$ and $3 + \sqrt{6}$. Then *their* square roots are in \mathbb{F}_{p^4} , so $\alpha = \sqrt{2 + \sqrt{2}} + \sqrt{3 + \sqrt{6}}$ is in \mathbb{F}_{p^4} . All the other roots of f can be expressed as polynomials in α (see Appendix C.2), so they are all in \mathbb{F}_{p^4} . Hence, the minimal polynomial of any root of f (over \mathbb{F}_p) divides $x^{p^4} - x$ by a standard characterization of finite fields. It follows that f factors over \mathbb{F}_p as a product of irreducible polynomials, each of degree 1, 2, or 4.

Suppose g is some irreducible factor of f . The ideal $(g(x))$ contains $(f(x))$, so there exists a ring homomorphism σ from $\mathbb{F}_p[x]/(f(x))$ to $\mathbb{F}_p[\alpha]/(g(x))$. Note that $\mathbb{F}_p[x]/(g(x))$ is a field because $g(x)$ is irreducible over \mathbb{F}_p . Also, σ fixes \mathbb{F}_p , so we obtain

$$\text{Per}(\sigma(G_1) \cdots \sigma(G_m)) = 2^a 3^b \Delta_C$$

as an equation over the field $\mathbb{F}_p[\alpha]/(g(x))$. For each i , $\sigma(G_i)$ is orthogonal in $\mathbb{F}_p[\alpha]/(g(x))$ as well:

$$\sigma(G_i)\sigma(G_i)^T = \sigma(G_i G_i^T) = \sigma(I) = I.$$

It follows that $M := \sigma(G_1) \cdots \sigma(G_m)$ is orthogonal.

Depending on the degree of g , the field $\mathbb{F}_p[\alpha]/(g(x))$ is isomorphic to \mathbb{F}_p , \mathbb{F}_{p^2} , or \mathbb{F}_{p^4} . But \mathbb{F}_{p^4} contains \mathbb{F}_p and \mathbb{F}_{p^2} , so M can be lifted to a matrix over \mathbb{F}_{p^4} . Given the permanent of M in \mathbb{F}_{p^4} , we can easily solve for Δ_C , so this completes the proof of Theorem 10.

Theorem 10 shows that for any prime $p \neq 2, 3$ there is some finite field of characteristic p where computing permanents (of orthogonal matrices) is hard. In particular, $p = 2$ and $p = 3$ are the only cases where the permanent of an orthogonal matrix is easy to compute in *every* finite field of characteristic p , assuming the polynomial hierarchy does not collapse. We will now show that there are primes p for which this problem is hard in *any* field of characteristic p , by showing that it is hard to compute in \mathbb{F}_p (which is contained in every other field of characteristic p).

Theorem 12. *For all primes p that split completely in $\mathbb{Q}(\alpha)$, computing the permanent of an orthogonal matrix over \mathbb{F}_p is Mod_pP -complete. This is a sequence of primes with density $\frac{1}{16}$ beginning*

191, 239, 241, 337, 383, 433, 673, 863, 911, 1103, 1151, 1249, 1583, 1871, 1873, 2017, ...

Proof. Recall that in the proof of Theorem 10, if g is an irreducible factor of f , then the result applies over the field $\mathbb{F}_p[x]/(g(x)) \cong \mathbb{F}_{p^{\deg g}}$. We show that g is degree at most 4, but in special cases this can be improved. In particular, we want g to be degree 1 (i.e., a linear factor) for our orthogonal matrix to be over \mathbb{F}_p .

First, observe that $\mathbb{Q}(\alpha)$ is a Galois extension of \mathbb{Q} . That is, every root of the minimal polynomial for α is in $\mathbb{Q}(\alpha)$. See Appendix C.2 for details.

We apply Chebotarev's density theorem [29], which says that if K is a finite Galois extension of \mathbb{Q} of degree n , then the density of primes which split completely in K is $\frac{1}{n}$. We take $K = \mathbb{Q}(\alpha)$, a degree 16 extension of \mathbb{Q} .

For our purposes, a prime p splits completely if and only if f (the minimal polynomial for α) factors into distinct linear terms modulo p . Conversely, since $\mathbb{Q}(\alpha)$ is a Galois extension, f will split into equal degree factors. Hence, if any factor is linear, then *all* the factors are linear. There are only finitely many primes which *ramify*, meaning that some factor is repeated; for primes which do not ramify, there is a linear factor if and only if the prime splits completely.

According to Chebotarev’s theorem, about 1/16 of all primes split completely and yield the desired hardness result. We verified the list of primes given in the theorem computationally. Also note that for about 3/16 of all primes, we can prove a hardness result over \mathbb{F}_{p^2} but not \mathbb{F}_p . \square

We leave open how hard it is to compute the permanent of an orthogonal matrix modulo the remaining $\frac{15}{16}$ of all primes. Other linear optical gadgets can be used for CSIGN instead of V , resulting in different field extensions where different primes split. For instance, there exists an orthogonal gadget for KLM’s NS_1 gate for which computing the permanent modulo 97 is hard (see Appendix F). However, it seems impossible to design linear optical gadgets that do not involve 2 or 3 photons at a time, in which case writing down $\varphi(O)$ requires $\sqrt{2}$ and $\sqrt{3}$. By quadratic reciprocity, these square roots only exist if $p \equiv \pm 1 \pmod{24}$ (i.e., for about a quarter of all primes), so the remaining primes may require some other technique.

5 Expanding Permanent Hardness

In this section, we try to fill in some of the remaining landscape of matrix permanents. In particular, we will focus on the permanents of positive semidefinite (PSD) matrices and their connection to boson sampling. We will conclude by listing some matrix variants and their accompanying permanent complexities, many of which are simple consequences of the reduction in Section 3.

5.1 Positive Semidefinite Matrix Permanents

Permanents of PSD matrices have recently become relevant to the expanding theory of boson sampling [21]. Namely, permanents of PSD matrices describe the output probabilities of a boson sampling experiment in which the input is a tensor product of thermal states. That is, each of the m modes of the system starts in a state of the form

$$\rho_i = (1 - \tau_i) \sum_{n=0}^{\infty} \tau_i^n |n\rangle\langle n|$$

where $\tau_i = \langle n_i \rangle / (\langle n_i \rangle + 1)$ and $\langle n_i \rangle$ is average number of photons one observes when measuring ρ_i . In particular, notice that $\tau_i \geq 0$.

Let U be a unitary matrix representing a linear optical network. Let D be the diagonal matrix with τ_1, \dots, τ_m along the diagonal, and $A := UDU^\dagger$. Since $\tau_i \geq 0$ for all i , A is PSD. We can calculate the probability of detecting one photon in each mode:¹²

$$\langle 1|^{\otimes m} \left(\varphi(U) \left(\bigotimes_{i=1}^m \rho_i \right) \varphi(U)^\dagger \right) |1\rangle^{\otimes m} = \frac{\text{Per}(A)}{\prod_{i=1}^m (1 + \langle n_i \rangle)}.$$

One might then reasonably ask, “how hard is it to compute such probabilities?” The following theorem answers that question in the exact case.

Theorem 13. *The permanent of a positive-definite matrix in $\mathbb{Z}^{n \times n}$ is #P-hard. This implies #P-hardness for the larger class of positive semidefinite matrices.*

¹²A similar formula arises for detecting 1 photon in each of k distinct modes and 0 photons in the remaining $m - k$ modes.

Proof. It is well-known that the permanent of a 0-1 matrix is #P-hard [31]. Therefore, let $B \in \{0, 1\}^{n \times n}$ and consider the matrix

$$\Lambda_B = \begin{pmatrix} 0 & B \\ B^T & 0 \end{pmatrix}$$

Since $\text{Per}(B) \geq 0$, we have $\text{Per}(B) = \sqrt{\text{Per}(\Lambda_B)}$. Also observe that $\Lambda_B^T = \Lambda_B$, so Λ_B is Hermitian, that is, diagonalizable with real eigenvalues. Furthermore, since B is a 0-1 matrix, its spectral radius is at most $2n$. Defining $\Lambda_B(x) := \Lambda_B + xI$, we see that $\Lambda_B(x)$ is positive-definite for all $x > 2n$.

Notice now that $\text{Per}(\Lambda_B(x))$ is a degree- $2n$ polynomial in x . Therefore, given an oracle that calculates the permanent of a positive-definite matrix, we can interpolate a monic polynomial through the points $x = 2n + 1, 2n + 2, \dots, 4n$ to recover the polynomial $\text{Per}(\Lambda_B(x))$. Since $\text{Per}(\Lambda_B(0)) = \text{Per}(\Lambda_B)$, the permanent of a positive-definite matrix under *Turing reductions* is #P-hard.

We now only have left to prove that the above reduction can be condensed into a single call to the positive-definite matrix permanent oracle. Since the matrix B is a 0-1 matrix, the polynomial $\text{Per}(\Lambda_B(x))$ has positive integer coefficients, the largest of which is at most $(2n)!$. Therefore, if $x > (2n)!$, then we can simply read out each of the coefficients of $\text{Per}(\Lambda_B(x))$ with a single oracle call. Clearly, this requires at most a polynomial increase in the bit length of the integers used in the reduction. \square

Theorem 13 implies that there is some linear optical experiment one can perform with thermal input states for which calculating the exact success probability is computationally difficult. We would like to say that this also precludes an efficient classical sampling algorithm (unless PH collapses), as is done in work by Aaronson and Arkhipov [3] and Bremner, Jozsa, Shepherd [8]. Unfortunately, those arguments rely on the fact that even finding an *approximation* to their output probabilities is difficult, but the following theorem heavily suggests that such a result cannot exist.

Theorem 14 (Rahimi-Keshari, Lund, Ralph [21]). *There exists an efficient classical sampling algorithm for Boson Sampling with thermal input states. Furthermore, multiplicatively approximating the permanent of a PSD matrix is in the class FBPP^{NP}.*

Intuitively, such an algorithm exists because it is possible to write the permanent of a PSD matrix as an integral¹³ of a nonnegative function, on which we can use Stockmeyer's approximate counting algorithm [24]. Such a representation as a sum of positive terms also implies that the permanent of a PSD matrix is nonnegative, which we can also prove using the convexity arguments in Appendix E (see Corollary 31).

Notice that this also justifies our use of techniques distinct from the linear optical approach. Suppose we can encode the answer to a GapP-hard problem into the permanent of a PSD matrix as we do with real orthogonal matrices, then multiplicatively approximating the permanent of a PSD matrix would also be GapP-hard under Turing reductions (see Theorem 26 in Appendix D). On the other hand, Theorem 14 says that such a multiplicative approximation *does* exist, so

$$\text{PH} \subseteq \text{P}^{\text{GapP}} \subseteq \text{BPP}^{\text{NP}} \subseteq \Sigma_3^{\text{P}}.$$

¹³Suppose we have PSD matrix $A = CC^\dagger$ where $C = \{c_{i,j}\}$. Then the permanent of A can be expressed as the following expected value over complex Gaussians:

$$\text{Per}(A) = \mathbb{E}_{x \in \mathcal{G}_C(0,1)^n} \left[\prod_{i=1}^n \left| \sum_{j=1}^n c_{i,j} x_j \right|^2 \right].$$

Therefore, either such a reduction does not exist or the polynomial hierarchy collapses to the third level.

5.2 More Permanent Consequences of the Main Result

In this section, we try to give a sense in which our proof for the hardness of the permanent for real orthogonal matrices leads to new hardness results for many classes of matrices. The structure of this section is as follows: we will first *restrict* as much as possible the class of matrices for which the permanent is #P-hard; we will then observe that the permanent for any larger class of matrices must also be hard, which will show hardness for many natural classes of matrices.

We call matrix A an *involution* if $A = A^{-1}$.

Theorem 15. *Let A be a real orthogonal involution with $\text{Per}(A) \geq 0$. The permanent of A is #P-hard.*

Proof. Let $C : \{0, 1\}^n \rightarrow \{0, 1\}$ be a Boolean function for which we want to calculate Δ_C . We will construct a new circuit $C' : \{0, 1\}^{n+1} \rightarrow \{0, 1\}$ such that for $x \in \{0, 1\}^n$ and $b \in \{0, 1\}$ we have $C'(x, b) = C(x) \vee b$. It is not hard to see then that $\Delta_{C'} = \Delta_C + 2^n$. Importantly, this implies that $\Delta_{C'} \geq 0$.

Now let us leverage the reduction in Theorem 7 to build a real orthogonal matrix B such that $\text{Per}(B) \propto \Delta_{C'}$. As in the proof of Theorem 13, let

$$\Lambda_B = \begin{pmatrix} 0 & B \\ B^T & 0 \end{pmatrix}.$$

Since $\Delta_{C'} \geq 0$, we have $\text{Per}(B) \geq 0$, which implies that $\text{Per}(B) = \sqrt{\text{Per}(\Lambda_B)}$. However, since B is orthogonal, we have that $\Lambda_B^2 = I$, so Λ_B is an involution. Furthermore, $\Lambda_B = \Lambda_B^T$, so Λ_B is a real orthogonal matrix. Therefore, the permanent of real orthogonal involutions is #P-hard. \square

We call a matrix A *special* if $\det(A) = 1$. Furthermore, a matrix A is *symplectic* if $A^T \Omega A = \Omega$ where $\Omega = \begin{pmatrix} 0 & I_n \\ -I_n & 0 \end{pmatrix}$. We strengthen Theorem 15 to provide the smallest class of matrices for which we know the permanent is #P-hard.

Theorem 16. *Let A be a real special orthogonal symplectic involution with $\text{Per}(A) \geq 0$. The permanent of A is #P-hard.*

Proof. Let B be a real orthogonal involution, and let I_n be the $n \times n$ identity matrix. Consider the matrix

$$I_2 \otimes B = \begin{pmatrix} B & 0 \\ 0 & B \end{pmatrix}.$$

Notice that

$$\det(I_2 \otimes B) = \det(B)^2 = \det(B^2) = \det(I_n) = 1,$$

where we use that $B^2 = I_n$ is an involution for the third equality. Therefore, $I \otimes B$ is special. It is also easy to verify that $I_2 \otimes B$ is real orthogonal symplectic involution. Assuming $\text{Per}(B) \geq 0$, we have $\text{Per}(B) = \sqrt{\text{Per}(I_2 \otimes B)}$. Combining the above with Theorem 15, we get that the permanent of real special orthogonal involutions is #P-hard. \square

Since the set of $n \times n$ real special orthogonal matrices form a group $\text{SO}(n, \mathbb{R})$, we immediately get #P-hardness for all the matrix groups containing it.

Corollary 17. *The permanent of an $n \times n$ matrix A in any of the classical Lie groups over the complex numbers is $\#P$ -hard. That is, it is hard for the following matrix groups:*

- General linear:** $A \in \text{GL}(n)$ iff $\det(A) \neq 0$
- Special linear:** $A \in \text{SL}(n)$ iff $\det(A) = 1$
- Orthogonal:** $A \in \text{O}(n)$ iff $AA^T = I_n$
- Special orthogonal:** $A \in \text{SO}(n)$ iff $AA^T = I_n$ and $\det(A) = 1$
- Unitary:** $A \in \text{U}(n)$ iff $AA^\dagger = I_n$
- Special unitary:** $A \in \text{SU}(n)$ iff $AA^\dagger = I_n$ and $\det(A) = 1$
- Symplectic:** $A \in \text{Sp}(2n)$ iff $A^T \Omega A = \Omega$ where $\Omega = \begin{pmatrix} 0 & I_n \\ -I_n & 0 \end{pmatrix}$

Proof. Since $\text{SO}(n, \mathbb{R})$ is a subgroup of all the stated Lie groups besides the symplectic group $\text{Sp}(2n)$, their permanents are $\#P$ -hard by Theorem 16. Theorem 16 handles the symplectic case separately. \square

6 Open Problems

This paper gives many new classes of matrices for which the permanent is hard. Nevertheless, there exist classes of matrices which have unknown permanent complexity, and proving $\#P$ -hardness or otherwise remains a central open problem. For instance, is computing the permanent of an orthogonal matrix modulo a prime p hard for all $p \neq 2, 3$? Notice that our result only gives $\text{Mod}_p P$ -hardness for 1/16th of all primes.

Another interesting open question about permanents concerns the complexity of multiplicatively approximating permanents of PSD matrices. Although we show the exact version of this problem to be $\#P$ -hard in this paper, we know that an FBPP^{NP} algorithm exists [21]. Could this problem actually just be in P ? Is there any more insight to be gained by viewing PSD permanents as probabilities of certain boson sampling experiments? For instance, Chakhmakhchyan, Cerf, and Garcia-Patron [10] have recently detailed conditions on the eigenvalues of a PSD matrix for which a linear optical sampling algorithm gives a better *additive* approximation to the permanent than the classic approximation algorithm of Gurvits [14].

Finally, we ask what other sorts of classical problems might benefit from linear optical (and more generally, quantum) techniques. Indeed, linear optics seemed almost indispensable in proving that the permanent of a real orthogonal matrix is $\#P$ -hard. To take another example of such a quantum-classical connection, Kuperberg [19] shows that computing certain values of the Jones polynomial to high accuracy is $\#P$ -hard using $\text{PostBQP} = \text{PP}$, a well-known result of Aaronson [1]. For a more thorough treatment of this topic, see the survey on quantum proofs for classical theorems of Drucker and de Wolf [11].

7 Acknowledgments

We would like to thank Scott Aaronson for posing the question which led to this paper and for his comments on this paper. We would also like to thank Rio LaVigne for some useful number theory pointers.

References

- [1] S. Aaronson. Quantum computing, postselection, and probabilistic polynomial-time. *Proc. Roy. Soc. London*, A461(2063):3473–3482, 2005. quant-ph/0412187.
- [2] S. Aaronson. A linear-optical proof that the permanent is $\#P$ -hard. *Proc. Roy. Soc. London*, A467(2088):3393–3405, 2011. arXiv:1109.1674.
- [3] S. Aaronson and A. Arkhipov. The computational complexity of linear optics. *Theory of Computing*, 9(4):143–252, 2013. Conference version in Proceedings of ACM STOC’2011. ECCC TR10-170, arXiv:1011.3245.
- [4] S. Aaronson, D. Grier, and L. Schaeffer. The classification of reversible bit operations. *arXiv preprint arXiv:1504.05155*, 2015.
- [5] L. M. Adleman, J. DeMarrais, and M.-D. A. Huang. Quantum computability. *SIAM Journal on Computing*, 26(5):1524–1540, 1997.
- [6] A. Ben-Dor and S. Halevi. Zero-one permanent is $\#P$ -complete, a simpler proof. In *ISTCS*, pages 108–117, 1993.
- [7] S. J. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Information processing letters*, 18(3):147–150, 1984.
- [8] M. Bremner, R. Jozsa, and D. Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proc. Roy. Soc. London*, A467(2126):459–472, 2010. arXiv:1005.1407.
- [9] E. R. Caianiello. On quantum field theory, 1: explicit solution of Dyson’s equation in electrodynamics without use of Feynman graphs. *Nuovo Cimento*, 10:1634–1652, 1953.
- [10] L. Chakhmakhchyan, N. J. Cerf, and R. Garcia-Patron. A quantum-inspired algorithm for estimating the permanent of positive semidefinite matrices. *arXiv preprint arXiv:1609.02416*, 2016.
- [11] A. Drucker and R. de Wolf. Quantum proofs for classical theorems. *Theory of Computing Graduate Surveys*, (2):1–54, 2011. arXiv:0910.3376, ECCC TR03-048.
- [12] S. Fenner, L. Fortnow, and S. Kurtz. Gap-definable counting classes. *Journal of Computer and System Sciences*, 48(1):116 – 148, 1994.
- [13] S. Fenner, F. Green, S. Homer, and R. Pruim. Quantum NP is hard for PH. In *Proceedings of 6th Italian Conference on theoretical Computer Science*, pages 241–252. Citeseer, 1998.
- [14] L. Gurvits. On the complexity of mixed discriminants and related problems. In Joanna Jedrzejowicz and Andrzej Szepietowski, editors, *Mathematical Foundations of Computer Science 2005, 30th International Symposium, MFCS 2005, Gdansk, Poland, August 29 - September 2, 2005, Proceedings*, volume 3618 of *Lecture Notes in Computer Science*, pages 447–458. Springer, 2005.
- [15] C. K. Hong, Z. Y. Ou, and L. Mandel. Measurement of subpicosecond time intervals between two photons by interference. *Phys. Rev. Lett.*, 59(18):2044–2046, 1987.

- [16] E. Knill. Quantum gates using linear optics and postselection. *Physical Review A*, 66(5), 2002.
- [17] E. Knill, R. Laflamme, and G. J. Milburn. A scheme for efficient quantum computation with linear optics. *Nature*, 409:46–52, 2001. See also quant-ph/0006088.
- [18] G. Kogan. Computing permanents over fields of characteristic 3: Where and why it becomes difficult. In *Foundations of Computer Science, 1996. Proceedings., 37th Annual Symposium on*, pages 108–114. IEEE, 1996.
- [19] G. Kuperberg. How hard is it to approximate the Jones polynomial? *Theory of Computing*, 11(6):183–219, 2015.
- [20] M. Nielsen and I. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [21] S. Rahimi-Keshari, A. P. Lund, and T. C. Ralph. What can quantum optics say about computational complexity theory? *Physical review letters*, 114(6):060501, 2015.
- [22] Terry Rudolph. Simple encoding of a quantum circuit amplitude as a matrix permanent. *Physical Review A*, 80(5):054302, 2009.
- [23] Y. Shi. Both Toffoli and controlled-NOT need little help to do universal quantum computation. *Quantum Information and Computation*, 3(1):84–92, 2002. quant-ph/0205115.
- [24] L. J. Stockmeyer. The complexity of approximate counting. In *Proc. ACM STOC*, pages 118–126, 1983.
- [25] S. Toda. PP is as hard as the polynomial-time hierarchy. *SIAM J. Comput.*, 20(5):865–877, 1991.
- [26] S. Toda and M. Ogiwara. Counting classes are at least as hard as the polynomial-time hierarchy. *SIAM Journal on Computing*, 21(2):316–328, 1992.
- [27] T. Toffoli. Reversible computing. In *Proc. Intl. Colloquium on Automata, Languages, and Programming (ICALP)*, pages 632–644. Springer, 1980.
- [28] L. Troyansky and N. Tishby. Permanent uncertainty: On the quantum evaluation of the determinant and the permanent of a matrix. In *Proceedings of PhysComp*, 1996.
- [29] N. Tschebotareff. Die bestimmung der dichtigkeit einer menge von primzahlen, welche zu einer gegebenen substitutionsklasse gehören. *Mathematische Annalen*, 95(1):191–228, 1926.
- [30] L. G. Valiant. Completeness classes in algebra. In *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing, STOC '79*, pages 249–261, New York, NY, USA, 1979. ACM.
- [31] L. G. Valiant. The complexity of computing the permanent. *Theoretical Comput. Sci.*, 8(2):189–201, 1979.

A Counting Classes

Let us introduce the complexity classes we use in this paper. Note that the permanent is a function, so computing it is a function problem. Hence, we will sometimes need the class FP to stand in for P when we are talking about function problems.

Definition 18. *FP is the class of functions computable by deterministic Turing machines in polynomial time.*

Of course, computing the permanent is, in general, thought to be intractable (i.e., not in FP). We use a variety of different classes to capture the difficulty of computing the permanent (depending on the kind of matrix, underlying field, etc.), but the most important class is #P:

Definition 19. *#P is the class of function problems of the form “compute the number of accepting paths of a polynomial-time non-deterministic Turing machine.” For example, given a classical circuit of NAND gates as input, the problem of computing the number of satisfying assignments is in #P (and indeed, is #P-complete).*

Since #P is a class of function problems (more specifically, counting problems), we often consider $P^{\#P}$ to compare #P to decision classes. Observe that $P^{\#P} = P^{PP}$ since, on the one hand, the #P oracle can count paths to simulate PP, and on the other hand, we can use the PP oracle to binary search (on the number of accepting paths) to count exactly. We add that $P^{\#P} \subseteq PSPACE$ is an upper bound for #P, and Toda’s theorem [25] gives $PH \subseteq P^{\#P}$.

Fenner, Fortnow, and Kurtz [12] define a very closely related class, GapP, which is also relevant to us.

Definition 20. *GapP is the class of function problems of the form “compute the number of accepting paths minus the number of rejecting paths of a polynomial-time non-deterministic Turing machine.”*

We have $GapP \supseteq \#P$ since we can take a #P problem (manifest as a non-deterministic Turing machine) and at the end of each rejecting path, add a non-deterministic branch which accepts in one half and rejects in the other. In the other direction, any GapP problem can be solved with at most two calls to a #P oracle (one for accepting paths, one for rejecting), and a subtraction. Hence, for most of our results we neglect the difference and stated #P-hardness.

Nonetheless, GapP and #P are different. For one, functions in #P are non-negative (and integral) by definition, whereas functions in GapP can take negative values. The distinction is also important in the context of approximation; Stockmeyer’s approximate counting gives a multiplicative approximation to any #P problem in BPP^{NP} , whereas it is known that multiplicative approximation to a GapP-hard problem remains GapP-hard under Turing reductions (see Theorem 26).

One cannot even get very bad multiplicative approximations to GapP-hard problems. Even the worst multiplicative approximation will distinguish zero from non-zero outputs, and this problem is captured by the class $C=P$, defined below.

Definition 21. *$C=P$ is the class of decision problems of solvable by a non-deterministic polynomial-time machine which accepts if it has the same number of accepting paths as rejecting paths.*

A good upper bound for $C=P$ is simply PP. This is easily seen once we have the following theorem.

Theorem 22. *Suppose $f_1, f_2 \in \Sigma^* \rightarrow \mathbb{Z}$ are functions computable in GapP. Then $f_1 + f_2$, $-f_1$, and $f_1 f_2$ are computable in GapP.*

Proof. Let M_1 and M_2 be non-deterministic machines witnessing $f_1 \in \text{GapP}$ and $f_2 \in \text{GapP}$ respectively. Then the machines for $f_1 + f_2$, $-f_1$, and $f_1 f_2$ are defined as follows.

1. For $f_1 + f_2$, non-deterministically branch at the start, then run M_1 in one branch and M_2 in the other.
2. For $-f_1$, take the complement of M_1 . That is, make every accepting path reject, and make every rejecting path accept.
3. For $f_1 f_2$, run M_1 to completion, then run M_2 to completion (in every branch of M_1). Accept if the two machines produce the same outcome, otherwise reject.

The last construction may require some explanation. Let a_1, a_2 be the number of accepting paths of M_1 and M_2 respectively, and similarly let b_1, b_2 be the numbers of rejecting paths. Then there are $a_1 a_2 + b_1 b_2$ accepting paths for the new machine and $a_1 b_2 + a_2 b_1$ rejecting paths, so as a GapP machine it computes

$$a_1 a_2 - a_1 b_2 - a_2 b_1 + b_1 b_2 = (a_1 - b_1)(a_2 - b_2) = f_1(x)f_2(x).$$

□

Theorem 22 implies that $\text{C=P} \subseteq \text{PP}$ because we can square and negate the gap. In other words, we can find a machine such that the gap is always negative (i.e., strictly less than half of all paths accept) unless the original machine had gap zero, in which case the gap is still zero (or, WLOG, very slightly positive). It is also worth noting that coC=P is known to equal NQP , by a result of Fenner et al. [13].

Definition 23. *The class NQP contains decision problems solvable by a polynomial-time quantum Turing machine (or, equivalently, a uniform, polynomial-size family of quantum circuits) where we accept if there is any nonzero amplitude on the accept state at the end of the computation.*

Quantum classes with exact conditions on the amplitudes (e.g., NQP or EQP) tend to be very sensitive to the gate set, or QTM transition amplitudes allowed. Adleman, Demarrais, and Huang [5] are careful to define NQP for the case where the transition amplitudes are algebraic and real.

Finally, we specify computational hardness for our finite field problems using a mod k decision version of $\#P$.

Definition 24. *For any integer $k \geq 2$, let $\text{Mod}_k P$ be the class of decision problems solvable by a polynomial time non-deterministic machine which rejects if the number of accepting paths is divisible by k , and accepts otherwise. In the special case $k = 2$, $\text{Mod}_k P$ is also known as “parity P ”, and denoted $\oplus P$.*

Clearly $P^{\#P}$ is an upper bound for $\text{Mod}_k P$. We are finally ready to state the main hardness result for these counting classes, namely, the celebrated theorem of Toda [25] and a subsequent generalization by Toda and Ogiwara [26]. There are many important consequences of Toda’s work, but we only require the following formulation.

Theorem 25 (Toda’s Theorem [25, 26]). *Let A be one of the counting classes $\text{Mod}_k P$, C=P , $\#P$, PP , or GapP . Then $\text{PH} \subseteq \text{BPP}^A$.*

This means in particular that, if a problem is hard for any of these classes, then there is no efficient algorithm for the problem unless PH collapses.

B Real Construction of Toffoli (Proof of Lemma 5)

In this appendix we prove Lemma 5 from Section 3. Let us first define R_θ as the rotation by θ about the Y -axis. That is, $R_\theta = \cos(\theta/2)I - i \sin(\theta/2)Y$ where Y is the Pauli σ_Y matrix. For our purposes, we only require the following two matrices:

$$R_{\pi/4} = \frac{1}{2} \begin{pmatrix} \sqrt{2 + \sqrt{2}} & -\sqrt{2 - \sqrt{2}} \\ \sqrt{2 - \sqrt{2}} & \sqrt{2 + \sqrt{2}} \end{pmatrix} \quad R_\pi = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$$

Let us now recall the statement of the lemma:

Lemma 5. *There exists a circuit of CSIGN, Hadamard, and $R_{\pi/4}$ gates which implements a Toffoli gate exactly.*

Proof. We construct the Toffoli gate from the CSIGN, Hadamard, and $R_{\pi/4}$ gates in three steps:

1. **Construct a controlled-controlled- R_π gate (CC- R_π) from CSIGN and $R_{\pi/4}$ gates.** CC- R_π is a three-qubit gate that applies R_π to the third qubit if the first two qubits are in the state $|11\rangle$. Notice that CC- R_π is already a kind of ‘‘poor man’s’’ Toffoli gate. If it were not for the minus sign in the R_π gate, we would be done. The construction is given in Figure 3. Observe that if either of the two control qubits is zero, then any CNOT gate controlled by that qubit can be ignored. The remaining gates will clearly cancel to the identity. Furthermore, if the two control qubits are in the state $|11\rangle$, then on the last qubit, we apply the operation $XR_{\pi/4}^{-1}XR_{\pi/4}XR_{\pi/4}^{-1}XR_{\pi/4}$. Since $XR_{\pi/4}^{-1}X = R_{\pi/4}$,

$$XR_{\pi/4}^{-1}XR_{\pi/4}XR_{\pi/4}^{-1}XR_{\pi/4} = R_{\pi/4}^4 = R_\pi.$$

Notice that this construction uses CNOT gates, but observe that a CNOT is a CSIGN gate conjugated by the Hadamard gate:

$$(I \otimes H) \text{CSIGN}(I \otimes H) = \text{CNOT}.$$

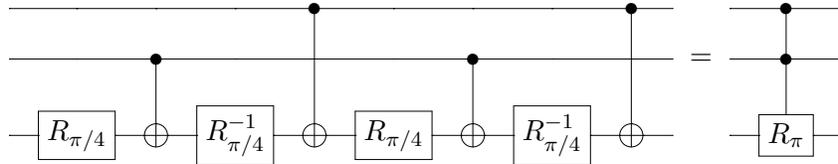


Figure 3: Generating CC- R_π from the CNOT and $R_{\pi/4}$ gates.

2. **Construct a non-affine classical reversible gate from CSIGN and CC- R_π gates.** By classical, we simply mean that the gate maps each computational basis state to another computational basis state (i.e., states of the form $|x\rangle$ for $x \in \{0, 1\}^n$). If this transformation is non-affine, then it suffices to generate Toffoli (perhaps with some additional ancilla qubits) by Aaronson et al. [4]. The construction is shown in Figure 4.
3. **Use the non-affine gate to generate Toffoli.** We give an explicit construction in Figure 5. Notice that the fourth qubit is an ancillary qubit starting in the $|0\rangle$ state.¹⁴

□

¹⁴Indeed, this ancillary qubit is necessary because the non-affine gate in Figure 4 is an even permutation and the Toffoli gate is an odd permutation on three bits.

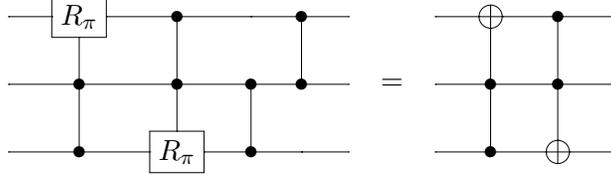


Figure 4: Generating non-affine classical gate from CC- R_π and CSIGN.

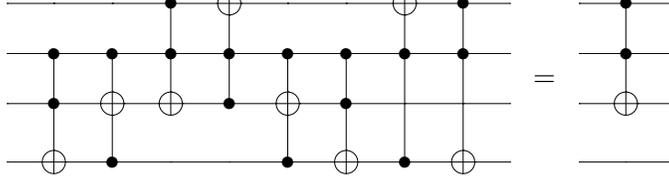


Figure 5: Generating Toffoli gate from non-affine gate in Figure 4.

C Gadget Details

As discussed above in Section 3 and Section 4, our results on orthogonal matrices depend on a collection of gadgets. In the real orthogonal setting (Section 3), each gadget is a real orthogonal matrix with algebraic entries, and all entries have clear, compact expressions in terms of radicals. However, in Section 4, we wish to reuse the same gadgets over finite fields, and radicals are no longer the best representation.

Instead, we will show that our (real) gadget matrices have entries in $\mathbb{Q}(\alpha)$, the algebraic field extension of the rational numbers by α , where $\alpha = \sqrt{2 + \sqrt{2}} + \sqrt{3 + \sqrt{6}} \approx 4.182173283$ is the largest real root of irreducible polynomial

$$f(x) = x^{16} - 40x^{14} + 572x^{12} - 3736x^{10} + 11782x^8 - 17816x^6 + 11324x^4 - 1832x^2 + 1.$$

More specifically, we will write every entry as a polynomial in α , with rational coefficients and degree less than 16.

This is a cumbersome representation for hand calculation, but there are some advantages. First, it eliminates any ambiguity about, for instance, which square root of 2 to use in a finite field. Second, we can check the various conditions our gadgets need to satisfy in the field $\mathbb{Q}(\alpha)$, and then argue that the verification generalizes to $\mathbb{F}_p(\alpha)$, with a few caveats. So, without further ado, we present polynomials for a set of reals which generate all the entries of our gadgets.

C.1 Gadget entries

Since $\pm \frac{1}{\sqrt{2}}$ are the only entries in D , our decoder gadget, we show how to express those entries as polynomials in α .

$$\frac{1}{\sqrt{2}} = \frac{1}{11776} (\alpha^{14} - 53\alpha^{12} + 1077\alpha^{10} - 10561\alpha^8 + 51555\alpha^6 - 115791\alpha^4 + 95207\alpha^2 - 8379),$$

For our encoder gadget E , we also must also express $\frac{1}{\sqrt{3}}$ as an element in $\mathbb{Q}(\alpha)$. Note that $\frac{1}{\sqrt{6}}$ can be obtained as $\frac{1}{\sqrt{2}} \cdot \frac{1}{\sqrt{3}}$.

$$\frac{1}{\sqrt{3}} = \frac{1}{11776} (\alpha^{14} - 53\alpha^{12} + 1077\alpha^{10} - 10561\alpha^8 + 51555\alpha^6 - 115791\alpha^4 + 95207\alpha^2 - 8379).$$

Showing that the entries of the $R_{\pi/4}$ gate are in $\mathbb{Q}(\alpha)$ requires the following:

$$\begin{aligned}\sqrt{2 + \sqrt{2}} &= \frac{1}{5888} (-123\alpha^{15} + 4932\alpha^{13} - 70785\alpha^{11} + 464494\alpha^9 \\ &\quad - 1470141\alpha^7 + 2209176\alpha^5 - 1357287\alpha^3 + 193302\alpha) \\ \sqrt{2 - \sqrt{2}} &= \frac{1}{5888} (216\alpha^{15} - 8711\alpha^{13} + 126234\alpha^{11} - 841629\alpha^9 \\ &\quad + 2733428\alpha^7 - 4270353\alpha^5 + 2799098\alpha^3 - 466411\alpha)\end{aligned}$$

Finally, we have the V gate. We already have the $\frac{1}{3\sqrt{2}}$ in front, and the various multiples of $\sqrt{2}$ inside, so we just need $\sqrt{3 \pm \sqrt{6}}$ and $\sqrt{6 \pm 2\sqrt{6}}$. These are related by a factor of $\sqrt{2}$, so it suffices to give $\sqrt{3 \pm \sqrt{6}}$.

$$\begin{aligned}\sqrt{3 + \sqrt{6}} &= \frac{1}{5888} (123\alpha^{15} - 4932\alpha^{13} + 70785\alpha^{11} - 464494\alpha^9 \\ &\quad + 1470141\alpha^7 - 2209176\alpha^5 + 1357287\alpha^3 - 187414\alpha) \\ \sqrt{3 - \sqrt{6}} &= \frac{1}{256} (15\alpha^{15} - 598\alpha^{13} + 8505\alpha^{11} - 55084\alpha^9 \\ &\quad + 171665\alpha^7 - 256518\alpha^5 + 161671\alpha^3 - 25624\alpha)\end{aligned}$$

The numbers above, combined with $\frac{1}{2}$ and $\frac{1}{3}$, generate all the entries of our real orthogonal gadgets. Note that the denominators in front of the polynomials above (e.g., 11776, 5888, 256, 3, etc.) all divide $35328 = 2^9 \cdot 3 \cdot 23$. In other words, this representation is a bad choice for fields of characteristic 2, 3, or 23 because, in those cases, division by 35328 is division by 0. Aside from this restriction, the representation is well-defined for any field containing some root α of the polynomial p .

We should not be surprised that the representation fails for fields of characteristic 2 or 3 because our matrices contain, for instance, the entries $\frac{1}{3}$ and $\frac{1}{\sqrt{2}}$. We also know the permanent of an orthogonal matrix is easy to compute in fields of characteristic 2 or 3, so it is actually no surprise to find this obstacle to our hardness proof.

On the other hand, we can find no explanation for the requirement $p \neq 23$; it appears to be a quirk of the algebraic number α . In fact, a different choice fails for different primes. Consider $\beta \approx 5.596386846$, the largest real root of

$$\begin{aligned}x^{16} - 56x^{14} - 32x^{13} + 1084x^{12} + 960x^{11} - 9224x^{10} - 8928x^9 + 37702x^8 + \\ 33920x^7 - 73736x^6 - 53216x^5 + 63932x^4 + 23488x^3 - 21560x^2 + 3808x - 191.\end{aligned}$$

This appendix is long enough without doing all the same steps for β , so let us claim without proof that $\mathbb{Q}(\beta) = \mathbb{Q}(\alpha)$. Furthermore, when we represent the matrix entries as polynomials in β (we omit the details), the denominators prohibit the use of this representation for fields of characteristic 2, 3, 191, and 3313, but *not* 23. Hence, for all primes p other than 2 or 3, there is some representation that works for that prime.

C.2 Galois Extension

We need $\mathbb{Q}(\alpha)$ to be a Galois extension to apply Chebotarev's theorem, which we use to prove Theorem 12. Another helpful consequence is that if α is in some field, then all the roots of f are also in the field since they can be expressed as polynomials in α .

The most direct way to prove $\mathbb{Q}(\alpha)$ is a Galois extension is to write all 16 roots of f in terms of α . Since f is an even polynomial, half of the roots are just the negatives of the other half, so we restrict our attention to the 8 positive roots.

Root	Polynomial
0.0234	$\frac{1}{5888} (-129\alpha^{15} + 5043\alpha^{13} - 69381\alpha^{11} + 425303\alpha^9 - 1214867\alpha^7 + 1629561\alpha^5 - 919335\alpha^3 + 122941\alpha)$
0.4866	$\frac{1}{2944} (123\alpha^{15} - 4932\alpha^{13} + 70785\alpha^{11} - 464494\alpha^9 + 1470141\alpha^7 - 2209176\alpha^5 + 1357287\alpha^3 - 190358\alpha)$
1.1057	$\frac{1}{2944} (-234\alpha^{15} + 9343\alpha^{13} - 133200\alpha^{11} + 865713\alpha^9 - 2709218\alpha^7 + 4054545\alpha^5 - 2537860\alpha^3 + 391327\alpha)$
1.5073	$\frac{1}{5888} (561\alpha^{15} - 22465\alpha^{13} + 321849\alpha^{11} - 2108561\alpha^9 + 6681723\alpha^7 - 10170267\alpha^5 + 6517531\alpha^3 - 1055763\alpha)$
1.5690	$\frac{1}{5888} (-93\alpha^{15} + 3779\alpha^{13} - 55449\alpha^{11} + 377135\alpha^9 - 1263287\alpha^7 + 2061177\alpha^5 - 1441811\alpha^3 + 278997\alpha)$
2.5897	$\frac{1}{2944} (111\alpha^{15} - 4411\alpha^{13} + 62415\alpha^{11} - 401219\alpha^9 + 1239077\alpha^7 - 1845369\alpha^5 + 1180573\alpha^3 - 198025\alpha)$
3.0997	$\frac{1}{5888} (339\alpha^{15} - 13643\alpha^{13} + 197019\alpha^{11} - 1306123\alpha^9 + 4203569\alpha^7 - 6479529\alpha^5 + 4156385\alpha^3 - 653825\alpha)$
4.1821	α

D Approximation

Much like in Aaronson’s paper [2], our hardness reductions for *exactly* computing the permanent lead naturally to hardness of approximation results as well. Approximation results comes in two flavors: additive and multiplicative. For example, Gurvits’ algorithm [14] approximates the permanent of a matrix A up to $\pm\epsilon\|A\|^n$ additive error. We will focus strictly on multiplicative approximation. That is, the result of the approximation should be between $\frac{1}{k}\text{Per}(A)$ and $k\text{Per}(A)$ for some k .

We give approximation results only for real orthogonal matrices since it is unclear how to even define multiplicative approximation in a finite field. All of our results follow from the fact that we actually prove GapP -hardness (since we compute the gap, Δ_C , rather than just the number of satisfying assignments). None of the results use anything specific to permanents; they are all GapP folklore, but we state them as permanent results for clarity.

Theorem 26. *Suppose A is an oracle that approximates the permanent of a real orthogonal matrix to any multiplicative factor. In other words, A is an oracle for the sign (zero, positive, or negative) of the permanent. Then $\text{GapP} \subseteq \text{FP}^A$.*

Proof. We give an FP^A algorithm for computing Δ_C for a classical circuit C . Since this problem is GapP -hard, we get $\text{GapP} \subseteq \text{FP}^A$.

By earlier results, we can construct a real orthogonal matrix with permanent proportional to Δ_C . Then we can apply the oracle to compute the sign of the permanent, and hence the sign of Δ_C . This is helpful, but we can do better.

Recall that we can add or subtract two GapP functions (see Appendix A), so for any integer k , we can construct a circuit C_k such that $\Delta_{C_k} = \Delta_C - k$. Then we can apply A to give us the sign of Δ_{C_k} , or equivalently, compare Δ_C to k . In other words, we can use A to binary search for the value of Δ_C , which we know to be an integer in the range -2^n and 2^n . \square

Recall that C=P is the class of decision problems of solvable by a non-deterministic polynomial-time machine which accepts if it has the same number of accepting paths as rejecting paths. By Toda’s theorem, $\text{PH} \subseteq \text{BPP}^{\text{C=P}}$.

Theorem 27. *Suppose A is an oracle that approximates the absolute value of permanent of a real orthogonal matrix to any multiplicative factor. That is, A tells us whether the permanent is zero. Then $\text{P}^{\text{C=P}} \subseteq \text{P}^A$.*

Proof. The problem of computing whether $\Delta_C = 0$ for a classical circuit C is C=P -hard. But clearly we can construct a real, orthogonal matrix from the circuit with permanent proportional to Δ_C , and then apply A to determine if the permanent is zero, and hence whether Δ_C is zero. Therefore $\text{P}^{\text{C=P}} \subseteq \text{P}^A$. \square

Finally, we show that even a very poor approximation to the absolute value of the permanent still allows us to calculate the exact value of the permanent via a boosting argument.

Theorem 28. *Suppose A is an oracle that approximates the absolute value of the permanent of an $n \times n$ real orthogonal matrix to within a $2^{n^{1-\varepsilon}}$ factor for some $\varepsilon > 0$. Then $\text{GapP} \subseteq \text{FP}^A$.*

Proof. We give an FP^A algorithm for computing Δ_C of a classical circuit. Since this problem is GapP -hard, we get $\text{GapP} \subseteq \text{FP}^A$.

As in Theorem 26, we can construct a circuit C_k such that $\Delta_{C_k} = \Delta_C - k$ for any integer k . By applying oracle A to the real orthogonal matrix corresponding to C_k , we can get a multiplicative estimate for $|\Delta_C - k|$. Let us assume for the moment that A gives a multiplicative approximation to within a factor of 2, and improve this to $2^{n^{1-\varepsilon}}$ later.

Suppose we are given an interval $[a, b]$ guaranteed to contain Δ_C . For instance, Δ_C is initially in $[-2^n, 2^n]$. Apply A to find an estimate for $\Delta_{C_a} = \Delta_C - a$. Suppose the approximation we get is x^* . Then we have

$$a + \frac{1}{2}x^* \leq \Delta_C \leq a + 2x^*.$$

So Δ_C is in the interval $[a + \frac{1}{2}x^*, a + 2x^*] \cap [a, b]$. One can show that this interval is longest when $a + 2x^* = b$, where it has length $\frac{3}{4}(b - a)$. Since the interval length decreases by a constant factor each step, we only need $O(n)$ steps to shrink it from $[-2^n, 2^n]$ to length < 1 , and determine Δ_C .

Finally, suppose we are given an oracle which gives an approximation to within a multiplicative factor $2^{n^{1-\varepsilon}}$. Theorem 22 in Appendix A lets us construct a circuit C^m (not to be confused with C_k) such that $\Delta_{C^m} = (\Delta_C)^m$. The circuit is essentially m copies of C , so we can only afford to do this for k polynomial in the size of C , otherwise our algorithm is too slow.

The point of C^m is that a factor β approximation to Δ_{C^m} gives a factor $\beta^{1/m}$ approximation of Δ_C by taking m th roots. This is excellent for reducing a constant approximation factor, but when β grows with n , we must account for the fact that the size of C^m grows with n as well. In particular, the size of C^m scales with m , and the dimension of the matrix in our construction scales linearly with m as well.

So, for our algorithm to succeed, we need $\beta(nm)^{1/m} \leq 2$ or

$$\beta(nm) \leq 2^m$$

for m a polynomial in n . Suppose we can afford $m = n^c$ copies of C . Then we succeed when $\beta(n^{1+c}) \leq 2^{n^c}$, or

$$\beta(n) \leq 2^{n^{1-\frac{1}{c+1}}}.$$

Within the scope of polynomial time algorithms, we can make $\frac{1}{c+1}$ less than any ε , and thereby handle any $2^{n^{1-\varepsilon}}$ approximation factor. \square

The core ideas in both Theorem 26 and 28 were already noticed by Aaronson [2], but we give slightly better error bounds for the latter theorem.

E Permanent Support via Linear Optics

In this appendix, we use the linear optical machinery to give a connection between the eigenvalues of a matrix and its permanent.

Theorem 29. *Suppose $A \in \mathbb{C}^{n \times n}$ is a normal matrix, i.e., $AA^\dagger = A^\dagger A$. Let $\Lambda \subseteq \mathbb{C}$ be the set of eigenvalues of A , and let $C = \{\lambda_1 \cdots \lambda_n : \lambda_1, \dots, \lambda_n \in \Lambda\} \subseteq \mathbb{C}$. That is, C is the set of products of n (possibly repeated) eigenvalues of A . Then $\text{Per}(A)$ is in the convex hull of C .*

Proof. Let us interpret A as a linear optical matrix, ignoring the fact that it might not be unitary. We know from the theory in Section 2 that

$$\text{Per}(A) = \langle 1 \cdots 1 | \varphi(A) | 1 \cdots 1 \rangle$$

where $|1 \cdots 1\rangle$ is the Fock state with one photon in every mode.

Now diagonalize A as UDU^\dagger , and use the fact that φ is a homomorphism.¹⁵

$$\begin{aligned} \text{Per}(A) &= \langle 1 \cdots 1 | \varphi(U) \varphi(D) \varphi(U)^\dagger | 1 \cdots 1 \rangle \\ &= \langle \psi | \varphi(D) | \psi \rangle \\ &= \text{Tr}(|\psi\rangle\langle\psi| \varphi(D)) \end{aligned}$$

where $|\psi\rangle = \varphi(U)^\dagger |1 \cdots 1\rangle$.

Observe that since D is a diagonal matrix (of the eigenvalues), $\varphi(D)$ will also be a diagonal matrix. Furthermore, each diagonal entry of $\varphi(D)$ is the product of n diagonal entries of D , i.e., an element of C .

In Section 2, we note that φ preserves unitarity, so $\varphi(U)^\dagger$ is a unitary matrix. It follows that

$$1 = \langle \psi | \psi \rangle = \text{Tr}(|\psi\rangle\langle\psi|).$$

We know the diagonal entries of the matrix $|\psi\rangle\langle\psi|$ are real, and by the above, they sum to 1. When we multiply $|\psi\rangle\langle\psi|$ by $\varphi(D)$, we scale the columns by diagonal entries of $\varphi(D)$ (which, recall, are elements of C). Hence, $\text{Per}(A) = \text{Tr}(|\psi\rangle\langle\psi| \varphi(U))$ is a convex combination of entries of C . \square

For certain types of matrices, the convex hull of products of eigenvalues is a geometrically meaningful region. In particular, we get easy corollaries (some of which were already known) about unitary and positive semidefinite matrices.

Corollary 30. *Suppose U is a unitary matrix. Then $|\text{Per}(U)|^2$ is at most 1.*

Corollary 31. *Let M be a positive semidefinite matrix. Then $\text{Per}(M)$ is real and non-negative. Furthermore, increasing the eigenvalues will increase the permanent. That is, if $M = UDU^\dagger$ and $M' = UD'U^\dagger$ where U is unitary, D, D' are diagonal, and $D \leq D'$ entrywise, then $\text{Per} M \leq \text{Per} M'$.*

Proof. The first part, $\text{Per}(M)$ is real and non-negative, follows directly from the theorem. When we increase the eigenvalues, we increase the diagonal of $\varphi(D)$, and therefore cannot decrease the permanent. \square

As far as the authors know, the latter part of Corollary 31, which connects the eigenvalues of a PSD matrix to its permanent, is a new result.

¹⁵Although not mentioned explicitly, unitarity is never used in the proof that φ is a homomorphism in Lemma 2 of Aaronson's paper [2].

F Orthogonal Matrices mod 97 are #P-hard via NS₁-approach

It is natural to ask whether Theorem 12 can be extended to more primes, or all primes. In other words, is there some prime $p \neq 2, 3$ such that it is easy to compute the permanent modulo p , even though computing the permanent over \mathbb{F}_{p^4} is hard? In this appendix, we present a different construction for CSIGN gates (in fact, the construction originally used by KLM) which works in \mathbb{F}_{97} , where the V gate does not. We conclude that there is at least one more prime, namely $p = 97$, where the permanent is hard.

The original KLM construction builds an CSIGN gate from what they call an NS₁ gate, instead of directly using a V gate. Logically, the NS₁ gate acts on one mode and does nothing to 0 or 1 photon, but flips the sign for 2 photons. The construction of CSIGN from NS₁ is shown in Figure 6. If $|1, 1\rangle$ is the input state, the Hadamard gate turns it into a linear combination of $|2, 0\rangle$ and $|0, 2\rangle$, which then change phase by the NS₁ gate, and get recombined into $-|1, 1\rangle$ by the Hadamard gate. Otherwise, there are not enough photons for the NS₁ gates to do anything, and the Hadamard gates cancel, so the gate does nothing (as a CSIGN should).

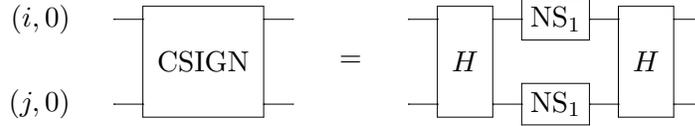


Figure 6: Generating CSIGN from H and NS₁ [17].

It turns out it is impossible to construct an NS₁ gate without at least two postselected modes, so the KLM NS₁ is a three mode gate where the last two modes start and end (via postselection) in state $|0, 1\rangle$. Unfortunately, the KLM NS₁ gate postselects on a mode having zero photons, which is undesirable for our application. Therefore, we construct our own NS₁ gate shown below. It postselects on the last two modes being $|1, 1\rangle$ and has entirely real entries.

The gate is

$$\text{NS}_1 = \frac{1}{6} \begin{pmatrix} 6 - 18\gamma & -\sqrt{6}\sqrt{9\gamma - \sqrt{6 - 3\gamma} - 2} & -\sqrt{6}\sqrt{9\gamma + \sqrt{6 - 3\gamma} - 2} \\ -\sqrt{6}\sqrt{9\gamma - \sqrt{6 - 3\gamma} - 2} & 9\gamma + \sqrt{24 - 45\gamma} & -3\sqrt{2 - 4\gamma} \\ -\sqrt{6}\sqrt{9\gamma + \sqrt{6 - 3\gamma} - 2} & -3\sqrt{2 - 4\gamma} & 9\gamma - \sqrt{24 - 45\gamma} \end{pmatrix}$$

where $\gamma \triangleq \frac{1}{18}(\sqrt{33} + 3) \approx 0.4858090359$.

One can verify the following identities hold.

$$\begin{aligned} \langle 0, 1, 1 | \phi(\text{NS}_1) | 0, 1, 1 \rangle &= \gamma, \\ \langle 1, 1, 1 | \phi(\text{NS}_1) | 1, 1, 1 \rangle &= \gamma, \\ \langle 2, 1, 1 | \phi(\text{NS}_1) | 2, 1, 1 \rangle &= -\gamma. \end{aligned}$$

That is, with amplitude γ the postselection succeeds, and the three mode gate behaves like an NS₁ gate on the first mode.

The field extension containing this gate is of higher degree than $\mathbb{Q}(\alpha)$, so we have not computed it explicitly. If we proved the equivalent of Theorem 12 in that extension, we would expect the density to be worse. However, this construction of an CSIGN works for at least one prime where V does not, namely $p = 97$.