Clifford Algebras and Approximating the Permanent

[Extended Abstract]

Steve Chien[*]
Computer Science Division
Univ. of California at Berkeley
Soda Hall, Berkeley CA
94720-1776
schien@cs.berkeley.edu

Lars Rasmussen[†]
Digital Fountain, Inc.
39141 Civic Center Drive,
Suite 300
Fremont, CA 94538
lars@digitalfountain.com

Alistair Sinclair[‡]
Computer Science Division
Univ. of California at Berkeley
Soda Hall, Berkeley CA
94720-1776
sinclair@cs.berkeley.edu

ABSTRACT

We study approximation algorithms for the permanent of an $n \times n$ $(0,1)$ matrix A based on the following simple idea: obtain a random matrix B by replacing each 1-entry of A independently by $\pm e$, where e is a random basis element of a suitable algebra; then output $|\det(B)|^2$. This estimator is always unbiased, but it may have exponentially large variance. In our first main result we show that, if we take the algebra to be a *Clifford algebra* of dimension polynomial in n, then we get an estimator with small variance. Hence only a constant number of trials suffices to estimate the permanent to good accuracy. The idea of using Clifford algebras is a natural extension of earlier work by Godsil and Gutman, Karmarkar *et al.*, and Barvinok, who used the real numbers, complex numbers and quaternions respectively.

The above result implies that, in principle, this approach gives a *fully-polynomial randomized approximation scheme* for the permanent, provided $|\det(B)|^2$ can be efficiently computed in the Clifford algebras. Since these algebras are non-commutative it is not clear how to do this. However, our second main result shows how to compute in polynomial time an estimator with the same mean and variance over the 4-dimensional algebra (which is the quaternions, and is non-commutative); in addition to providing some hope that the computations can be performed in higher dimensions, this quaternion algorithm provides an exponential improvement in the variance over that of the 2-dimensional complex version studied by Karmarkar *et al.*

1. INTRODUCTION

The permanent of an $n \times n$ $(0,1)$ matrix $A = (a_{ij})$ is defined as

$$\mathrm{per}(A) = \sum_{\pi \in S_n} \prod_{i=1}^{n} a_{i,\pi i}.$$

Equivalently, $\mathrm{per}(A)$ counts the perfect matchings in the $(n+n)$-vertex bipartite graph whose adjacency matrix is A. Computing $\mathrm{per}(A)$ exactly is #P-complete, as was shown in Valiant's seminal 1979 paper [18]. The best we can hope for therefore is an efficient approximation.

The past decade or so has seen a surprising variety of approaches aimed at designing a polynomial time approximation algorithm for the permanent. These can be divided into (at least) four categories: elementary recursive algorithms [16]; reductions to determinants [5, 10, 6, 2, 3]; iterative balancing [13]; and Markov chain Monte Carlo [4, 7, 9, 11, 8]. All the approaches have yielded non-trivial results (at a minimum, fully polynomial approximation schemes for random matrices, or polynomial time approximation algorithms with approximation ratio c^n for a modest constant $c > 1$), and fascinating insights both into the problem and into the wider implications of the associated mathematical techniques. Recently, as reported in [8], the Markov chain Monte Carlo approach led to the first *fully polynomial randomized approximation scheme* for the permanent of an arbitrary $(0,1)$ matrix (and indeed of any matrix with non-negative entries). This is a randomized algorithm which takes as inputs A and a parameter $\varepsilon \in (0,1]$ and outputs a value that approximates $\mathrm{per}(A)$ within a factor $1 \pm \varepsilon$ with high probability; the running time is polynomial in n and $\frac{1}{\varepsilon}$.

In this paper we pursue another of the approaches mentioned above, namely reduction to determinants. We are motivated both by the intrinsic elegance of this approach, and by the fact that, if successful, it seems likely to lead to a more efficient algorithm. (The authors of [8] did not attempt to minimize the exponent in their polynomial running time; but even with fine tuning that algorithm is unlikely to be practical.)

The origins of the determinant approach go back to the following beautiful observation of Godsil and Gutman [5]. Let A be an $n \times n$ $(0,1)$ matrix, and let B be the matrix obtained by replacing each 1-entry of A independently by a uniform random element of $\{\pm 1\}$. Then the random variable $(\det(B))^2$ is an *unbiased estimator* of $\mathrm{per}(A)$, i.e., its

[*] Supported by an NSF Graduate Fellowship and NSF grants CCR-9820951 and CCR-0121555.

[†] Work done while at Computer Science Division, UC Berkeley. Supported by NSF grant CCR-9505448.

[‡] Supported in part by NSF grants CCR-9820951 and CCR-0121555 and DARPA cooperative agreement F30602-00-2-0601.

STOC'02, May 19-21, 2002, Montreal, Quebec, Canada.
Copyright 2002 ACM 1-58113-495-9/02/0005 ...$5.00.

expectation is exactly $\mathrm{per}(A)$. This is easy to verify using the facts that the terms in the expansions of permanent and determinant are identical up to sign, and that every cross term in the expansion of $\det(B)^2$ disappears because it contains an independent factor b_{ij} with $\mathrm{E}[b_{ij}] = 0$.

Unfortunately, however, the above estimator has in general a very large variance, so if we were to use it to estimate $\mathrm{per}(A)$ we would need to take the mean of exponentially many independent samples to get a good estimate with high probability. More precisely, given any unbiased estimator X_A of $\mathrm{per}(A)$, the number of samples needed to approximate $\mathrm{per}(A)$ within a factor $1 \pm \varepsilon$ with high probability is $\frac{\mathrm{const}}{\varepsilon^2} \frac{\mathrm{E}[X_A^2]}{\mathrm{E}[X_A]^2}$. We call the ratio $\frac{\mathrm{E}[X_A^2]}{\mathrm{E}[X_A]^2}$ of the second moment to the square of the expectation the *critical ratio* of the estimator.

Karmarkar *et al.* [10] showed that the critical ratio of Godsil and Gutman's estimator when run on any $n \times n$ $(0,1)$ matrix A is bounded above by $3^{n/2}$. More remarkably, they also showed that if each 1-entry of A is replaced not by $\{\pm 1\}$ but by a random element of $\{\pm 1, \pm i\}$ (where i is the complex square root of -1),* forming a matrix C, then the analogous estimator $|\det(C)|^2$ is still unbiased and the bound on the critical ratio drops to $2^{n/2}$. This is still exponential, but substantially smaller than the Godsil-Gutman version. In addition, Frieze and Jerrum [6] showed that the critical ratio of the Karmarkar *et al.* estimator is polynomially bounded with high probability for a *random* $(0,1)$ matrix A.

These ideas were pushed further by Barvinok [2] in a rather different framework. Instead of asking how much time is needed to compute a $(1 \pm \varepsilon)$ approximation of $\mathrm{per}(A)$, Barvinok asked how good an approximation can be obtained in polynomial time. Under this measure, he showed that the original Godsil-Gutman idea could also be improved by replacing each 1-entry of A by a continuous sample from a standard normal distribution; the resulting algorithm approximates $\mathrm{per}(A)$ within a factor of about $(3.57)^n$ with high probability in polynomial time. Moreover, the extension of Barvinok's algorithm to the complex numbers, analogous to that of Karmarkar *et al.*, provides an improvement of this approximation ratio to about $(1.79)^n$. Finally, Barvinok also showed that a further extension to the *quaternion* algebra (i.e., each 1-entry of A is replaced by a value $b_1 + b_2 i + b_3 j + b_4 k$, where i, j, k are Hamilton's quaternions and the b's are independent standard normal) again improves the approximation ratio to about $(1.32)^n$. In a subsequent paper [3] Barvinok proposes an extension of his techniques to higher dimensional algebras and conjectures that, for sufficiently high dimension, it may yield a polynomial time approximation algorithm for the permanent within ratio c^n for any desired constant $c > 1$. (Note that this is a much weaker requirement than that of a fully polynomial randomized approximation scheme.)

In this paper, we extend these ideas to higher dimensional algebras in the original approximation scheme framework, and demonstrate a more dramatic improvement than that conjectured by Barvinok. As observed by Barvinok, the sets $\{1\}$, $\{1, i\}$, $\{1, i, j, k\}$ are the basis elements of the first three *Clifford algebras*†, of dimensions 1, 2 and 4 respectively. For each $m \geq 1$, we define a permanent estimator based on the Clifford algebra CL_m of dimension 2^{m-1}. The estimator is analogous to that of Godsil-Gutman and Karmarkar *et al.*, and is very easy to describe: simply replace each 1-entry of A by an independent element chosen u.a.r. from $\{\pm e_1, \pm e_2, \ldots, \pm e_{2^{m-1}}\}$, where the e_i are the basis elements of CL_m, to obtain a matrix B; then output $|\det(B)|^2$. Note that $\det(B)$ is a value in CL_m; the norm-square function $|\cdot|^2$ is simply the sum of squares of the coefficients in the above basis. It is not hard to show that this estimator remains unbiased for all m.

Our first main result is that the critical ratio of the estimator decreases dramatically with the dimension. Specifically, we show:

Theorem A *Let $X_A = |\det(B)|^2$ be the value output by the above algorithm over CL_m, with $m = 4q + 2$. Then $\mathrm{E}[X_A] = \mathrm{per}(A)$ and the critical ratio $\frac{\mathrm{E}[X_A^2]}{\mathrm{E}[X_A]^2}$ for any $n \times n$ matrix A is bounded above by $(1 + \frac{1}{2^{2q}})^{n/2}$.*

An immediate corollary of this theorem is that, if we put $q = \lceil \frac{1}{2} \log_2 n \rceil$, then the critical ratio is bounded by a constant! — i.e., a constant number of trials suffice to get a good approximation of $\mathrm{per}(A)$. Moreover, to achieve this we need only work in the algebra CL_{4q+2} of dimension $2^{4q+1} = O(n^2)$, which is also polynomial in n. Thus, in principle, the approach yields a fully-polynomial randomized approximation scheme for the permanent.‡

The only catch is that our estimator requires the computation of $|\det(B)|^2$, where B is a matrix of basis elements of a high-dimensional algebra. The algebras CL_m are *not commutative* for $m \geq 3$ ($m = 1$ is the reals; $m = 2$ is the complex numbers; $m = 3$ is the quaternions), so standard polynomial time determinant computations break down. (In fact, it is known that computing general determinants in a non-commutative setting is computationally infeasible [15].) Nonetheless, we are able to overcome this obstacle at least for the first interesting case, namely the quaternions; for this algebra, our general analysis shows that the critical ratio is at most $(3/2)^{n/2}$. In our second main result, we show how to define a modified quaternion estimator closely related to the original one, but easily computable in polynomial time. Surprisingly, we show that this modified estimator has the same first and second moments as the original one, yielding:

Theorem B *There is a quaternion-based unbiased estimator for the permanent that is computable in polynomial time and has critical ratio at most $(3/2)^{n/2}$.*

Recall that the estimator of Karmarkar *et al.* has critical ratio $2^{n/2}$, so Theorem B gives a further significant exponential improvement. We leave as an intriguing open problem the question of whether the higher-dimensional estimators with small variance whose existence is guaranteed by Theorem A also have modified versions that are computable in polynomial time.

The following is a brief road-map of the paper. In section 2 we present the minimal set of facts about Clifford algebras

*Actually Karmarkar *et al.* used complex *cube* roots of unity, rather than fourth roots as stated here. We prefer the latter choice as it fits more naturally into our generalized framework. It is not hard to check that the use of kth roots for any $k \geq 3$ leads to essentially the same asymptotic behavior.

†More accurately, the *second* Clifford algebras. For definitions see the next section.

‡We have chosen the values $m \equiv 2 \bmod 4$ to allow the cleanest statement of Theorem A. In fact the critical ratio is monotonically decreasing with m, and our techniques allow a similar bound to be computed for any m.

required to understand our work. We go on in section 3 to define our generalized estimators based on the Clifford algebras CL_m, and show that they are always unbiased. The bulk of this section is devoted to the proof of the bound on the critical ratio of these estimators, as stated in Theorem A. The proof exploits the substantial group-theoretic structure underlying the Clifford algebras, and offers as a byproduct new insights into why the introduction of complex numbers by Karmarkar *et al.* improves on the initial Godsil-Gutman algorithm. In section 4 we define and analyze the modified estimator over the quaternion algebra, thus proving Theorem B. This analysis has a small amount in common with that of Barvinok's continuous quaternion version [2], but the two differ substantially in that we are analyzing the second moment while he was analyzing the tails. Finally, we note that due to space limitations some of the proof details are deferred to the full version of the paper.

2. CLIFFORD ALGEBRAS

In this section we cover the necessary fundamentals of Clifford algebras that we require for our estimators. There is a great deal of theory on Clifford algebras, but we will present only the minimal required background. For further reading we recommend, e.g., [12].

The Clifford algebras we will use are real algebras with basis elements of the form $u_{a_1 a_2 \ldots a_k}$, with $a_i \in [m]$ and $a_i < a_{i+1}$, together with $u_\epsilon = 1$. The multiplication rules are simple: for $i \neq j \in [m]$, $u_i u_j = u_{ij} = -u_j u_i$, and $u_i u_i = 1$. A general element of the Clifford algebra can thus be written as $h = \sum_S c_S u_S$, where c_S is real and S ranges over subsets of $[m]$.

We will restrict ourselves to the "second Clifford algebras," in which every basis element must have an even number of subscripts: e.g., u_{12}, u_{2467} etc. We denote the mth such Clifford algebra CL_m. Clearly the number of basis elements of CL_m is the number of even-cardinality subsets of $[m]$, which is 2^{m-1}. Of course this is also the dimension of the algebra over the reals.

The first few Clifford algebras are familiar enough. CL_1 has 1 as its only basis element and is just the real numbers. CL_2 has basis $\{1, u_{12}\}$, and is in fact the complex numbers with $i = u_{12}$. (Note that $u_{12}^2 = u_1 u_2 u_1 u_2 = -u_1^2 u_2^2 = -1$.) CL_3 has basis $\{1, u_{12}, u_{23}, u_{13}\}$ and is the quaternions, with $u_{23} = j$ and $u_{13} = k$. (The reader may check the familiar properties $i^2 = j^2 = k^2 = -1$ and $ij = k = -ji$.) Note that CL_3 (and hence CL_m for all $m \geq 3$) is not commutative; however, if two basis elements do not commute then their two products differ only up to a sign. The reader may consult the multiplication table for CL_4 in the Appendix.

Conjugation in CL_m is defined in the natural way. The conjugate of a basis element u_S, written $\overline{u_S}$, is its inverse, i.e., the (unique) element that satisfies $u_S \overline{u_S} = \overline{u_S} u_S = 1$. The conjugate of a general element $u = \sum_S c_S u_S$ is defined as $\overline{u} = \sum_S c_S \overline{u_S}$. Note that in general we cannot construct the inverse of u from $\overline{u}$, as $u\overline{u}$ may not be real; indeed, in CL_m for $m \geq 4$ not every element has an inverse.

The following useful observations can readily be verified from the above definitions:

1. A basis element u_S is *self-conjugate* (i.e., u_S is its own inverse, $u_S = \overline{u_S}$) if and only if S consists of $4k$ subscripts. If S consists of $4k+2$ subscripts, then $\overline{u_S} = -u_S$. Notice that for any S, $u_S^2 = \pm 1$.
2. Two basis elements commute if and only if the number of subscripts they share is even; i.e., $u_S u_{S'} = u_{S'} u_S$ if and only if $|S \cap S'| = 2k$. Otherwise, $u_S u_{S'} = -u_{S'} u_S$.
3. If u, u' are signed basis elements chosen independently and uniformly at random, then their product uu' is also a uniformly random signed basis element.

It will be convenient to associate with each Clifford algebra CL_m the group of its 2^m *signed* basis elements G_m. Each group element $\alpha \in G_m$ corresponds to either u_S or $-u_S$ for some $S \subseteq [m]$, and the group operation is simply multiplication as defined in CL_m. For $\alpha \in G_m$, we denote by u_α the corresponding signed basis element in CL_m. Thus an arbitrary element of CL_m can be written as $h = \sum_{\alpha \in G_m} c_\alpha u_\alpha$ for $c_\alpha \in \mathbb{R}^+$. We assume that c_α is non-zero for at most one of u_S and $-u_S$, so that this representation is unique.

Recall from the introduction that our permanent estimators are of the form $|\det(B)|^2$, where the entries of the matrix B, and therefore also $\det(B)$, lie in the Clifford algebra CL_m. Thus we need to define the norm-square $|u|^2$ for $u \in CL_m$. Generalizing from the reals, complex numbers and quaternions, we might try to use the definition $|u|^2 = u\overline{u}$. However, this is problematic in CL_m for $m \geq 4$ because $u\overline{u}$ is not in general real. Instead, we will define $|u|^2$ to be the *real part* of $u\overline{u}$; equivalently, if $u = \sum_S c_S u_S$, then $|u|^2 = \sum_S c_S^2$.

3. GENERAL CLIFFORD ALGEBRA ESTIMATORS

3.1 Definition and expectation

For each Clifford algebra CL_m, we can define a corresponding estimator for the permanent in the natural way: given a $(0,1)$ matrix A, replace each 1-entry of A with a signed basis element of CL_m chosen independently and uniformly at random to obtain a new matrix B; then compute $X_A = |\det(B)|^2$. Here $\det(B) = \sum_{\pi \in S_n} \mathrm{sgn}(\pi) \prod_{i=1}^n b_{i,\pi i}$, which is an element of the Clifford algebra CL_m; and $|\det(B)|^2$ is the real part of $\det(B)\,\overline{\det(B)}$. We prove first that this estimator is unbiased for all m. The proof is similar to the proofs for the low-dimensional versions of Godsil-Gutman ($m = 1$) and Karmarkar *et al.* ($m = 2$), but is complicated by the fact that $u\overline{u}$ is not necessarily real.

PROPOSITION 3.1. *In any Clifford algebra CL_m, we have* $\mathrm{E}[X_A] = \mathrm{per}(A)$.

Proof: We first introduce some notation. Given a permutation π, we define B_π to be $\prod_{i=1}^n b_{i,\pi i}$ and $\overline{B}_\pi$ to be $\prod_{i=n}^1 \overline{b}_{i,\pi i}$. Thus $B_\pi \overline{B}_\pi = \prod_{i=1}^n a_{i,\pi i}$. Further, given two permutations π_1 and π_2, we say that $\mathrm{R}(B, \pi_1, \pi_2) = 1$ if $B_{\pi_1} \overline{B}_{\pi_2}$ is real and 0 otherwise. Note that $\mathrm{R}(B, \pi, \pi) = 1$ for all π.

We can write the expectation of the estimator as $\mathrm{E}[X_A] = \sum_B \Pr(B) \sum_{\pi_1 \pi_2} \mathrm{sgn}(\pi_1 \pi_2) B_{\pi_1} \overline{B}_{\pi_2} \mathrm{R}(B, \pi_1, \pi_2)$ where the sum is over all possible choices of the random matrix B and $\Pr(B)$ is the probability of choosing B. We then proceed as follows:

$$\begin{aligned}
\mathrm{E}[X_A] &= \sum_{\pi_1} \sum_{B} \Pr(B) B_{\pi_1} \overline{B}_{\pi_1} \mathrm{R}(B, \pi_1, \pi_1) + \\
&\qquad\qquad \sum_{\pi_1 \neq \pi_2} \sum_{B} \Pr(B) B_{\pi_1} \overline{B}_{\pi_2} \mathrm{R}(B, \pi_1, \pi_2) \\
&= \sum_{\pi_1} \prod_{i} a_{i,\pi i} + \sum_{\pi_1 \neq \pi_2} \sum_{B} \Pr(B) B_{\pi_1} \overline{B}_{\pi_2} \mathrm{R}(B, \pi_1, \pi_2)
\end{aligned}$$

$$= \quad \mathrm{per}(A) + \sum_{\pi_1 \neq \pi_2} \sum_{B} \Pr(B) B_{\pi_1} \overline{B}_{\pi_2} \mathrm{R}(B, \pi_1, \pi_2).$$

To finish the proof, note that all choices of B are equally likely. When $\pi_1 \neq \pi_2$, let j be the smallest index such that $\pi_1 j \neq \pi_2 j$. Then $b_{j,\pi_1 j}$ is chosen independently of the other factors in the product $B_{\pi_1} \overline{B}_{\pi_2}$, and for each value u_S it takes on, it takes on $-u_S$ with equal probability; in each case the value of $R(B, \pi_1, \pi_2)$ is the same. Thus the sum on the right is 0 and $\mathrm{E}[X_A] = \mathrm{per}(A)$. □

3.2 The second moment: block diagonal case

Recall that the efficiency of the estimator X_A is governed by its critical ratio, $\frac{\mathrm{E}[X_A^2]}{\mathrm{E}[X_A]^2}$. Thus we need to compute the second moment, $\mathrm{E}[X_A^2]$. We first perform a detailed analysis for block diagonal matrices, and then in the next subsection use this to derive a bound for all matrices. Let $A_1 = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$. Then the block diagonal matrix A_n is the $2n \times 2n$ matrix with n copies of A_1 along its diagonal. From Proposition 3.1 we have $\mathrm{E}[X_{A_n}] = \mathrm{per}(A_n) = 2^n$. For convenience we define A_0 to be the 1×1 identity matrix; note that X_{A_0} is identically 1. We will study the distribution of X_{A_n} in the algebra CL_m as m varies.

The main result of this section is the following theorem:

Theorem 3.2. *For $n \geq 0$, let A_n be the above block diagonal matrix. Then in CL_{4q+2}, we have $\mathrm{E}[X_{A_n}^2] \leq [4(1 + \frac{1}{2^{2q+1}})]^n$, and thus the critical ratio $\frac{\mathrm{E}[X_{A_n}^2]}{\mathrm{E}[X_{A_n}]^2} \leq (1 + \frac{1}{2^{2q+1}})^n$.*

Let B_n denote the random matrix computed by the algorithm when run on A_n. Recall that the estimator X_{A_n} is defined as the real part of $\det(B_n)\,\overline{\det(B_n)}$. Since the product of two random signed basis elements in CL_m is again a random signed basis element, it is clear that $\det(B_1)$ has the same distribution as $a + b$, where a, b are independent random signed basis elements. Thus we may write (in distribution) $\det(B_n) = (a_1 + b_1)(a_2 + b_2) \cdots (a_n + b_n)$, where the a_i and b_i are randomly and independently chosen signed basis elements from CL_m. Therefore $\det(B_n)\,\overline{\det(B_n)} = (a_1 + b_1)(a_2 + b_2) \cdots (a_n + b_n)(\overline{a_n} + \overline{b_n}) \cdots (\overline{a_1} + \overline{b_1})$. Thus the distribution of $\det(B_{n+1})\,\overline{\det(B_{n+1})}$ is equivalent to that of $(a_1 + b_1)\det(B_n)\,\overline{\det(B_n)}(\overline{a_1} + \overline{b_1})$, which in turn is equivalent to

$$a(1 + c)\det(B_n)\,\overline{\det(B_n)}(1 + \overline{c})\overline{a}, \qquad (1)$$

where a, c are independent random signed basis elements. Since $\det(B_0)\overline{\det(B_0)}$ is always 1, expression (1) is valid for all $n \geq 0$.

The success of the Clifford algebra estimators can be explained by the algebraic restrictions on the behavior of X_{A_n}. The concrete ideas are contained in the following statement and its proof:

Lemma 3.3. *Suppose we are working in CL_m. Then the quantity $\det(B_n)\,\overline{\det(B_n)}$ is either zero or of the form $2^k \sum_{\alpha \in G} u_\alpha$, where (1) k is a non-negative integer; and (2) G is a* self-conjugate *subgroup of G_m (i.e., $u_\alpha^2 = 1$ for all $\alpha \in G$).*

Note that any self-conjugate subgroup G is necessarily abelian. Note also that the above representation assumes that G does not contain both α and $-\alpha$; all the self-conjugate subgroups in the sequel can easily be seen to have this property and we will assume it from now on. We will write $-G$ to denote the set of $\alpha \in G_m$ such that $-\alpha \in G$.

Proof of Lemma 3.3 (sketch): We proceed by induction on n. In the base case $n = 0$, we have $\det(B_0)\,\overline{\det(B_0)} = 1$, which can be written in the form $2^k \sum_{\alpha \in G} u_\alpha$ with $k = 0$ and G the trivial subgroup $\{1\}$. For the inductive step we need to look at $\det(B_{n+1})\,\overline{\det(B_{n+1})}$, which we recall can be written in the form given in expression (1). Applying the induction hypothesis to $\det(B_n)\,\overline{\det(B_n)}$ and expanding, this becomes

$$2^k a\Big(\sum_{\alpha \in G} u_\alpha + c(\sum_{\alpha \in G} u_\alpha)\overline{c} + c(\sum_{\alpha \in G} u_\alpha) + (\sum_{\alpha \in G} u_\alpha)\overline{c}\Big)\overline{a}. \quad (2)$$

Our task is therefore to show that, for an arbitrary self-conjugate subgroup G and signed basis elements a and c, the expression in (2) is either zero or can be written in the form $2^{k'} \sum_{\alpha \in G'} u_\alpha$ for some $k' \geq 0$ and self-conjugate subgroup G'. This follows from a fairly straightforward case analysis, whose structure we outline below. The proofs for each case are deferred to the Appendix.

The two main cases are case 1 (c commutes with all of G) and case 2 (c does not commute with all of G). We subdivide case 1 into three subcases according to whether or not c is self-conjugate (i.e., $c = \overline{c}$) and, if so, whether c lies in $G \cup -G$.

Case 1: c commutes with G.

CASE 1A: $c = \overline{c}$, $c \in G \cup -G$. Here if $c \in G$ then the outside coefficient 2^k quadruples to 2^{k+2}, and the subgroup G *transmutes* to another subgroup $G' = aG\overline{a}$ of the same size as G. Otherwise, if $c \in -G$, we obtain 0.

CASE 1B: $c = \overline{c}$, $c \notin G \cup -G$. Here the coefficient 2^k doubles to 2^{k+1}, while G *expands* to a subgroup $G' = a(G \cup cG)\overline{a}$ of twice the size of G.

CASE 1C: $c = -\overline{c}$. Here the coefficient doubles to 2^{k+1}, while G transmutes to $G' = aG\overline{a}$, a subgroup of the same size.

Case 2: c does not commute with G. As in case 1C, the coefficient doubles to 2^{k+1} while G transmutes to another subgroup G' of the same size.

This completes the sketch of the proof of Lemma 3.3. □

Example: We illustrate each of the above cases with a running example from CL_5, the 16-dimensional algebra whose subscripts are drawn from $\{1, \ldots, 5\}$. We start with the base case $\det(B_0)\,\overline{\det(B_0)} = 1$ (so $k = 0$ and G is the trivial subgroup $\{1\}$) and follow the evolution of $\det(B_n)\,\overline{\det(B_n)}$ for some particular sequence of choices of c. (We will fix $a = 1$ throughout for ease of computation.)

- $c = u_{1234}$ (case 1B). Expression (1) is $(1+u_{1234})(1)(1+u_{1234})$, which simplifies to $2(1 + u_{1234})$. The outside coefficient has doubled to 2 and G has expanded to the subgroup $\{1, u_{1234}\}$.

- $c = u_{1234}$ (case 1A). Here expresssion (1) is $2(1 + u_{1234})(1 + u_{1234})(1 + u_{1234})$, which simplifies to $8(1 + u_{1234})$. The outside coefficient has quadrupled to 8 and G has transmuted (to itself, because $a = 1$). (Note that if c had been $-u_{1234}$, we would have obtained 0.)

- $c = u_{23}$ (case 1C). Here expression (1) is $8(1+u_{23})(1+u_{1234})(1 - u_{23})$, which simplifies to $16(1 + u_{1234})$. The outside coefficient has doubled and G has transmuted (to itself).

- $c = u_{25}$ (case 2). Here expression (1) is $16(1+u_{25})(1+u_{1234})(1-u_{25})$, which simplifies to $32(1-u_{1345})$. The outside coefficient has doubled and G has transmuted to the subgroup $\{1, -u_{1345}\}$. □

The proof of Lemma 3.3 reveals a simple pattern to the behavior of the random variable $\det(B_n)\overline{\det(B_n)}$ that allows us to easily bound $\mathrm{E}[X_{A_n}^2]$. Note that since X_{A_n} is the real part of $\det(B_n)\overline{\det(B_n)}$, its value is just the outside coefficient 2^k in the representation in Lemma 3.3. From the above case analysis, we see that when n increases by 1 this coefficient either doubles or quadruples or drops to zero; and which of these outcomes occurs depends on the relationship between c and G and on whether c is self-conjugate. Since c is chosen u.a.r. from the signed basis elements, we can easily assign probabilities to each of these outcomes. This allows us to prove:

Lemma 3.4. *Let p be the maximum possible value of the ratio $2|G|/|G_m|$ over all self-conjugate subgroups G in G_m. Then in CL_m we have $\mathrm{E}[X_{A_n}^2] \le [4(1+p)]^n$, and thus the critical ratio $\frac{\mathrm{E}[X_{A_n}^2]}{\mathrm{E}[X_{A_n}]^2} \le (1+p)^n$.*

Proof: We again use induction on n. In the base case $n = 0$, we have $\mathrm{E}[X_{A_0}^2] = 1$. For the inductive step we examine the random variable $\det(B_{n+1})\overline{\det(B_{n+1})}$. Recall that, conditioned on the value of $\det(B_n)\overline{\det(B_n)}$, the distribution of this r.v. is as in (1) where a, c are independent random signed basis elements. From Lemma 3.3 we know that $\det(B_n)\overline{\det(B_n)}$ is either zero or of the form $2^k\sum_{\alpha\in G}u_\alpha$ for some k and G; thus the r.v. X_{A_n} has value zero or 2^k respectively. From the proof of Lemma 3.3 we see that the outside coefficient 2^k exactly doubles in all cases except case 1a. In this latter case it either quadruples (if $c \in G$) or becomes zero (if $c \in -G$). Plainly each of these outcomes occurs with probability $|G|/|G_m|$. Thus, conditioned on X_{A_n}, the distribution of $X_{A_{n+1}}$ is

$$\begin{cases} 0 & \text{with probability } |G|/|G_m|; \\ 2X_{A_n} & \text{with probability } 1-2|G|/|G_m|; \\ 4X_{A_n} & \text{with probability } |G|/|G_m|. \end{cases}$$

Since $|G|/|G_m| \le p/2$ by definition of p, we therefore have

$$\mathrm{E}[X_{A_{n+1}}^2] \le \left(16\tfrac{p}{2}+4(1-p)\right)\mathrm{E}[X_{A_n}^2] = 4(1+p)\mathrm{E}[X_{A_n}^2]. \quad (3)$$

This completes the proof by induction on n. □

Lemma 3.4 bounds the second moment in terms of $p = \max\frac{2|G|}{|G_m|}$, where the maximum is over all self-conjugate subgroups G in CL_m. The final ingredient is to show that p decreases rapidly as a function of m:

Lemma 3.5. *Let $m = 4q+2$. Then in CL_m, $p = \frac{1}{2^{2q+1}}$.*

Proof of Lemma 3.5 (sketch): We shall give a very simple argument that gives a slightly weaker bound, namely $p \le \frac{1}{2^q}$, and conveys the main idea. The additional factor of 2 in the exponent requires some slightly more detailed analysis (see the full proof in the Appendix).

Let G be a self-conjugate subgroup of G_m, and let H be the subgroup $G\cap G_{m-1}$. It is easy to check that either $H = G$ or $|H| = |G|/2$. This tells us that for any self-conjugate subgroup G of G_m, there is a subgroup of at least half its size in G_{m-1}. Hence p does not increase as m increases.

Now consider CL_{4q}. The basis element g that contains every index (e.g., $u_{12345678}$ in CL_8) is self-conjugate and commutes with every other basis element. This element (or its conjugate) must be contained in every maximal self-conjugate subgroup G of G_{4q}; otherwise $G\cup gG$ would be a larger such subgroup.

When we move to CL_{4q+1}, however, we see that any maximal self-conjugate subgroup G of G_{4q} cannot be augmented, since no element of $G_{4q+1}-G_{4q}$ commutes with g. Thus the size of the maximal subgroup in G_{4q+1} is unchanged from G_{4q}, so p decreases by a factor of 2. □

Putting Lemmas 3.4 and 3.5 together, we are done with the proof of Theorem 3.2 and hence the main business of this section. Theorem 3.2 and its proof contain the essential intuition about the behavior of the second moment on block diagonal matrices as m increases, and immediately imply that for $m = O(\log n)$, the critical ratio is bounded above by a constant. This proves Theorem A for the block-diagonal case. For technical reasons, in order to bootstrap this to a bound for general matrices we actually need to derive the exact form of $\mathrm{E}[X_{A_n}^2]$ as a sum of exponentials $\sum_i c_iE_i^n$, not just an upper bound $[4(1+\frac{1}{2^{2q+1}})]^n$ as in Theorem 3.2. This information is contained in the following theorem, whose proof follows from a slightly more refined analysis of the behavior of the subgroups than that used in the proof of Lemma 3.4 and may be found in the full version.

Theorem 3.6. *Let $m = 4q+2$. In CL_m, $\mathrm{E}[X_{A_n}^2] = c_1E_1^n + c_2E_2^n$, where $E_1 = 4(1+\frac{1}{2^{2q+1}})$, $0 < E_2 < E_1$, and c_1, c_2 are non-negative constants with $c_1+c_2 = 1$. Thus in particular $\mathrm{E}[X_{A_n}^2] \le E_1^n$.*

Remark: The reason we choose $m \equiv 2 \bmod 4$ is to allow the cleanest possible formulation of Lemma 3.5 and Theorem 3.6. However, the essential point is that a constant factor increase in the dimension (i.e., a constant additive increase in m) leads to a constant factor decrease in p (the relative size of the maximum subgroup). In fact, we have seen that p is monotonically decreasing with m, and a more detailed analysis shows that p decreases by a factor of 2 when $m \equiv 1, 2, 3$ or $5 \bmod 8$. □

3.3 The second moment: general case

In this section we extend our bound from the block diagonal case to prove the following result.

Theorem 3.7. *Let $m = 4q+2$. Then in CL_m, the critical ratio for an arbitrary $n\times n$ matrix A satisfies*

$$\frac{\mathrm{E}[X_A^2]}{\mathrm{E}[X_A]^2} \le \sum_{i=1}^{2} c_i\left(\frac{E_i-2}{2}\right)^{n/2} \le \left(\frac{E_1-2}{2}\right)^{n/2},$$

where c_1, c_2, E_1, E_2 are as in Theorem 3.6.

Substituting the value $E_1 = 4(1+\frac{1}{2^{2q+1}})$ immediately yields Theorem A in the Introduction.

Before embarking on the proof of Theorem 3.7, we introduce a useful graph-theoretic framework from [10]. Recall that we may view $\mathrm{per}(A)$ as the number of perfect matchings in a bipartite graph $\mathcal{B}(A)$ in which (i,j) is an edge if and only if $a_{ij} = 1$. We define $P(A)$ to be the set of permutations $\{\pi \in S_n : a_{i,\pi i} = 1 \,\forall i\}$, which correspond naturally to perfect matchings in $\mathcal{B}(A)$ (and we will blur this distinction). Thus $\mathrm{per}(A) = |P(A)|$. We need the following observations:

(i) Given two perfect matchings π_1 and π_2, let the graph G be their union $\pi_1 \cup \pi_2$. Then G is a disjoint union of even-length cycles and isolated edges. We will define $c(G)$ to be the number of cycles of G. Conversely, the cycle cover G can be described as the union of a pair of perfect matchings in $2^{c(G)}$ distinct ways. We write $\mathcal{G}(A)$ as the set of all such cycle covers. We therefore can write $\mathrm{per}(A)^2 = \sum_{G \in \mathcal{G}(A)} 2^{c(G)}$

(ii) Given $G, G' \in \mathcal{G}(A)$, we will say that $G' \subseteq G$ if all of the edges of G' are contained in G, or equivalently, if G' can be formed from G by collapsing some of the cycles of G. Thus there are $\binom{c(G)}{k} 2^k$ graphs $G' \subseteq G$ such that $c(G') = c(G) - k$.

(iii) Consider the union of four perfect matchings π_1, π_2, π_3, π_4. We will say that $\pi_1 \cup \pi_2 \cup \pi_3 \cup \pi_4$ is *even* if every edge in the union is covered an even number of times. In this case, $\pi_1 \cup \pi_2 \cup \pi_3 \cup \pi_4$ forms a cycle cover.

(iv) Consider any $G \in \mathcal{G}(A)$, and let $A(G)$ denote the adjacency matrix of G. Then the estimator $X_{A(G)}$ run on $A(G)$ has the same distribution as $X_{A_{c(G)}}$, the estimator on the block diagonal matrix with $c(G)$ blocks.

Proof of Theorem 3.7: Proceeding as in the proof of Proposition 3.1, we can write

$$\mathrm{E}[X_A^2] = \sum_B \Pr(B) \sum_{\pi_1\pi_2\pi_3\pi_4} \mathrm{sgn}(\pi_1\pi_2\pi_3\pi_4) \times$$
$$B_{\pi_1}\overline{B}_{\pi_2}B_{\pi_3}\overline{B}_{\pi_4}\mathrm{R}(B,\pi_1,\pi_2)\mathrm{R}(B,\pi_3,\pi_4) \qquad (4)$$

To simplify notation, define $B_{\pi_1\pi_2\pi_3\pi_4} = B_{\pi_1}\overline{B}_{\pi_2}B_{\pi_3}\overline{B}_{\pi_4}$ and write $\mathrm{RE}[B_{\pi_1\pi_2\pi_3\pi_4}]$ to denote the summation $\sum_B \Pr(B) B_{\pi_1\pi_2\pi_3\pi_4}\mathrm{R}(B,\pi_1,\pi_2)\mathrm{R}(B,\pi_3,\pi_4)$.

Our first observation is that $\mathrm{RE}[B_{\pi_1\pi_2\pi_3\pi_4}] = 0$ unless $\pi_1 \cup \pi_2 \cup \pi_3 \cup \pi_4$ is even. This follows because of the presence, in non-even cases, of an independent factor b in $B_{\pi_1\pi_2\pi_3\pi_4}$ that takes on values $\pm u_S$ with equal probability. Thus we may rewrite equation (4) as

$$\mathrm{E}[X_A^2] = \sum_{G \in \mathcal{G}(A)} \sum_{\pi_1 \cup \pi_2 \cup \pi_3 \cup \pi_4 = G} \mathrm{RE}[B_{\pi_1\pi_2\pi_3\pi_4}], \qquad (5)$$

where we can ignore $\mathrm{sgn}(\pi_1\pi_2\pi_3\pi_4)$ since it must be 1 when $\pi_1 \cup \pi_2 \cup \pi_3 \cup \pi_4$ is even.

We now prove, for any fixed $G \in \mathcal{G}(A)$, that

$$\sum_{\pi_1 \cup \pi_2 \cup \pi_3 \cup \pi_4 = G} \mathrm{RE}[B_{\pi_1\pi_2\pi_3\pi_4}] = \sum_{i=1}^{2} c_i (E_i - 2)^{c(G)}. \qquad (6)$$

This is done by induction on $c(G)$. The base case $c(G) = 0$ is verified by noticing that $\pi_1 = \pi_2 = \pi_3 = \pi_4$ must be the same permutation, so $B_{\pi_1\pi_2\pi_3\pi_4} = 1$, and both evaluations of $\mathrm{R}(\cdot)$ are also 1, so the left-hand side of (6) is $1 = \sum_i c_i (E_i - 2)^0$.

Now for any fixed G, let us define $A(G)$ as the $(0,1)$ matrix associated with G. Then

$$\mathrm{E}[X_{A(G)}^2] = \sum_{G' \in \mathcal{G}(A(G))} \sum_{\pi_1 \cup \pi_2 \cup \pi_3 \cup \pi_4 = G'} \mathrm{RE}[B_{\pi_1\pi_2\pi_3\pi_4}].$$

Since one possible instance of G' is G itself, we can rewrite this as

$$\begin{aligned}
&\sum_{\pi_1 \cup \pi_2 \cup \pi_3 \cup \pi_4 = G} \mathrm{RE}[B_{\pi_1\pi_2\pi_3\pi_4}] \\
&\quad = \mathrm{E}[X_{A(G)}^2] - \sum_{G' \subset G} \sum_{\pi_1 \cup \pi_2 \cup \pi_3 \cup \pi_4 = G'} \mathrm{RE}[B_{\pi_1\pi_2\pi_3\pi_4}] \\
&\quad = \sum_i c_i E_i^{c(G)} - \sum_{G' \subset G} \sum_{\pi_1 \cup \pi_2 \cup \pi_3 \cup \pi_4 = G'} \mathrm{RE}[B_{\pi_1\pi_2\pi_3\pi_4}] \\
&\quad = \sum_i c_i E_i^{c(G)} - \sum_{k=1}^{c(G)} \binom{c(G)}{k} 2^k \sum_i c_i (E_i - 2)^{c(G)-k} \\
&\quad = \sum_i c_i E_i^{c(G)} - \sum_i c_i (E_i^{c(G)} - (E_i - 2)^{c(G)}) \\
&\quad = \sum_i c_i (E_i - 2)^{c(G)}
\end{aligned}$$

In the second step here we have used the assumption of the theorem together with observation (iv) from earlier; in the third step we have used the induction hypothesis and observation (ii).

This completes the inductive proof of (6). Plugging the result into equation (5) gives

$$\mathrm{E}[X_A^2] = \sum_{G \in \mathcal{G}(A)} \sum_{i=1}^{2} c_i (E_i - 2)^{c(G)} \le \sum_{G \in \mathcal{G}(A)} (E_1 - 2)^{c(G)}.$$

Finally, combining this with the observation that $\mathrm{E}[X_A]^2 = \sum_{G \in \mathcal{G}(A)} 2^{c(G)}$ we obtain

$$\begin{aligned}
\frac{\mathrm{E}[X_A^2]}{\mathrm{E}[X_A]^2} &\le \frac{\sum_{G \in \mathcal{G}(A)} (E_1 - 2)^{c(G)}}{\sum_{G \in \mathcal{G}(A)} 2^{c(G)}} \\
&\le \max_{G \in \mathcal{G}(A)} \frac{(E_1 - 2)^{c(G)}}{2^{c(G)}} \\
&\le \left(\frac{E_1 - 2}{2}\right)^{n/2}.
\end{aligned}$$

This completes the proof of the theorem. □

We should note that our analysis includes the real- and complex-based estimators of Godsil-Gutman [5] and Karmarkar *et al.* [10] as special cases. In both cases the size of the maximum self-conjugate subgroup is 1, and so $E_1 = 8$ for $\mathbb{R}$ and $E_1 = 6$ for $\mathbb{C}$. Thus we see that the critical ratio is bounded by $3^{n/2}$ for $\mathbb{R}$ and $2^{n/2}$ for $\mathbb{C}$, which are the same as the bounds derived by less general methods in [10]. Our proof indicates how the group-theoretic structure of CL_m leads to the decrease with dimension of the critical ratio.

4. COMPUTING THE ESTIMATOR

We turn now to the question of implementing the estimators of the previous section. These estimators are defined in terms of the symbolic determinant of a matrix whose entries are basis elements of a high-dimensional Clifford algebra. Since such algebras are non-commutative, it is not clear how to perform such a computation in polynomial time; indeed, it is known that computing general determinants in a non-commutative setting is computationally infeasible [15].

Our goal in this final section is to show that this difficulty *can* be overcome at least in the first interesting case, namely the quaternion algebra $\mathbb{H} = CL_3$. (Recall that this algebra is non-commutative.) What we shall do is to construct a permanent estimator having the same flavor as that of the previous section, but which *is* efficiently computable; although the two estimators will not be equal, they will share the same first and second moments, so the performance guarantee of the previous section can actually be achieved in polynomial time.

4.1 A modified estimator over the quaternions

Our new estimator Y_A begins as before by replacing each 1-entry of A by a random element from $\mathbb{H}_8 = \{\pm 1, \pm i, \pm j, \pm k\}$, the signed basis elements of the quaternion algebra. Call the resulting random matrix H. Now, however, rather than working with the symbolic determinant $\det(H)$, we use its so-called *Dieudonné determinant*, defined as the result of performing a standard Gaussian elimination procedure on H as follows:

```
if $n = 1$ then return $\text{Gauss}(H) = h_{11}$
else if column $h_{\cdot 1} = 0$ then return $\text{Gauss}(H) = 0$
else if $h_{11} = 0$ then add any row $h_{i\cdot}$ with $h_{i1} \neq 0$ to row $h_{1\cdot}$
    for all $i > 1$ add row multiple $-h_{i1}h_{11}^{-1}h_{1\cdot}$ to row $h_{i\cdot}$
    return $\text{Gauss}(H) = h_{11}\text{Gauss}(H_{11})$
```

Note that $\text{Gauss}(H)$ is quaternion-valued; moreover, its value may depend on the row chosen in the third line. However, the classical theory of Dieudonné determinants (see, e.g., [1]) ensures that the norm-square, $|\text{Gauss}(H)|^2$, is well-defined (and indeed preserved under any sequence of row and column operations). Note that in the quaternions the norm-square is just $|h|^2 = h\overline{h}$ which is always real-valued, and hence inverses exist.

Evidently the estimator Y_A can be computed in $O(n^3)$ time. However, despite its resemblance to X_A of the previous section, it is not at all clear that it inherits the nice properties of that estimator. Indeed, it is not even clear that it is unbiased. Note in particular that $|\text{Gauss}(H)|^2$ and $|\det(H)|^2$ may differ considerably. For example, if we take $H = \begin{pmatrix} i & j \\ j & i \end{pmatrix}$ then it is easy to check that $|\text{Gauss}(H)|^2 = 4$ whereas $|\det(H)|^2 = 0$. However, we shall see presently that, when the non-zero entries of H are random quaternion basis elements, then these two quantities share the same first and second moments!

It will be convenient to note that $|\text{Gauss}(H)|^2$ can be written equivalently as a single complex determinant, known as the "reduced norm" of H. This is derived from the representation of the quaternions as 2×2 complex matrices as follows. For $h \in \mathbb{H}$, write h uniquely as $b + cj$, where $b, c \in \mathbb{C}$. Then h is represented by the matrix $\phi(h) = \begin{pmatrix} b & c \\ -\overline{c} & \overline{b} \end{pmatrix}$. Thus in particular the basis elements are represented as

$$\phi(1) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}; \qquad \phi(i) = \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix}; \qquad\qquad (7)$$

$$\phi(j) = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}; \qquad \phi(k) = \begin{pmatrix} 0 & i \\ i & 0 \end{pmatrix}.$$

Given an $n \times n$ quaternion matrix H, define a $2n \times 2n$ complex matrix $D = D_H$ by

$$d_{ij} = \left[\phi(h_{\lceil \frac{i}{2} \rceil, \lceil \frac{j}{2} \rceil})\right]_{(i \bmod 2),(j \bmod 2)}.$$

In words, D_H is formed by replacing each entry h_{ij} of H by its corresponding 2×2 complex matrix, and then erasing the boundaries between these matrices.

We define the *reduced norm* of H as $\det(D_H)$. (Since D_H is a complex matrix, this is well-defined and efficiently computable.) The following fact is easy to check (see, e.g., [2]):

PROPOSITION 4.1. *For any $n \times n$ quaternion matrix H, $|\text{Gauss}(H)|^2 = \det(D_H)$.*

Thus in what follows we may think of Y_A as being defined either as $|\text{Gauss}(H)|^2$ or as $\det(D_H)$. This flexibility will prove useful in our analysis.

Our goal is to show that the above permanent estimator Y_A is unbiased and has the same second moment as the quaternion version of the general estimator X_A derived in the previous section. Thus we will prove the following, which is exactly Theorem B of the Introduction.

THEOREM 4.2. *For any $n \times n$ $(0,1)$ matrix A, the Dieudonné determinant estimator Y_A satisfies*

$$\mathrm{E}[Y_A] = \mathrm{per}(A) \qquad \text{and} \qquad \frac{\mathrm{E}[Y_A^2]}{\mathrm{E}[Y_A]^2} \le \left(\tfrac{3}{2}\right)^{n/2}.$$

Our analysis will proceed along similar lines to that of the previous section, but the Dieudonné determinant will prove a little harder to work with than the symbolic determinant. In section 4.2 we will deal with the expectation, and in section 4.3 with the second moment.

4.2 Analysis of expectation

For the expectation, it will be useful to work with the reduced norm formulation of Y_A. Let $D = D_H \in \mathbb{C}^{2n \times 2n}$ be the reduced norm matrix computed by the algorithm. Then we have

$$\mathrm{E}[Y_A] = \mathrm{E}[\det(D)] = \mathrm{E}\Big[\sum_{\pi \in S_{2n}} \mathrm{sgn}(\pi) \prod_{i=1}^{2n} d_{i,\pi i}\Big] \stackrel{\text{def}}{=} \sum_{\pi \in S_{2n}} \mathrm{E}[D_\pi].$$

Clearly each entry d_{ij} of D depends on exactly one entry of H, namely $h_{\lceil \frac{i}{2} \rceil, \lceil \frac{j}{2} \rceil}$. Conversely, each entry h_{ij} of H determines four entries of D, namely $d_{2i-1,2j-1}$, $d_{2i-1,2j}$, $d_{2i,2j-1}$ and $d_{2i,2j}$. Moreover, by (7) we can view these four entries as b_{ij}, c_{ij}, $-\overline{c}_{ij}$ and $\overline{b}_{ij}$ respectively, where b_{ij}, c_{ij} are chosen randomly as follows: flip a fair coin. If Heads, choose b_{ij} u.a.r. from $\mathbb{C}_4 = \{\pm 1, \pm i\}$ and set $c_{ij} = 0$; if Tails, set $b_{ij} = 0$ and choose c_{ij} u.a.r. from $\mathbb{C}_4$.

Now let $\pi \in S_{2n}$. The factors $d_{i,\pi i}$ of D_π depend on a set $\pi|_n$ of between n and $2n$ entries of H, viz.

$$\begin{aligned} \pi|_n &= \{(\lceil \tfrac{i}{2} \rceil, \lceil \tfrac{\pi i}{2} \rceil) : i \in [2n]\} \\ &= \Big\{(k, \lceil \tfrac{\pi(2k-1)}{2} \rceil), (k, \lceil \tfrac{\pi(2k)}{2} \rceil) : k \in [n]\Big\}. \end{aligned}$$

The result is an immediate consequence of the following two claims:

CLAIM 1: Let $\pi \in S_{2n}$. Then $\mathrm{E}[D_\pi] \neq 0 \Rightarrow \pi|_n \in P(A)$.

CLAIM 2: Let $\sigma \in P(A)$. Then $\mathrm{E}[D_\sigma] \stackrel{\text{def}}{=} \sum_{\pi : \pi|_n = \sigma} \mathrm{E}[D_\pi] = 1$.

The intuition behind these claims is that the only permutations with nonzero expectation in D are those that correspond exactly to nonzero permutations in H, and each such permutation contributes an expected value of 1, thus yielding the permanent.

To prove Claim 1, consider $\pi \in S_{2n}$ with $\mathrm{E}[D_\pi] \neq 0$. Fix any odd $i = 2k-1 \in [2n]$, and assume $\pi i = 2l - 1$ is odd (the case of πi even is handled similarly). Thus $d_{i,\pi i} = b_{kl} \in \mathbb{C}_4$, where b_{kl} is the result of the random experiment described earlier. Since b_{kl} has a random sign, this factor cannot be independent of all other factors in D_π. But since the elements of H are all independent, the only other elements of D which are not independent of b_{kl} are $d_{i,\pi i+1} = c_{kl}$, $d_{i+1,\pi i} = -\overline{c}_{kl}$ and $d_{i+1,\pi i+1} = \overline{b}_{kl}$. And since π is a permutation, the only one of these that can be a factor of D_π is $d_{i+1,\pi i+1}$. Hence we must have $\pi(i+1) = \pi i + 1$. This in turn implies that $\lceil \frac{\pi i}{2} \rceil = \lceil \frac{\pi(i+1)}{2} \rceil$, so $\pi|_n$ contains only one entry in the kth row of H, namely $(k, \lceil \frac{\pi(2k-1)}{2} \rceil)$. Since i was arbitrary, we conclude that $\pi|_n \in S_n$. Clearly $\pi|_n$ cannot contain an index pair (k,l) with $a_{kl} = 0$, since otherwise D_π would be zero. Thus $\pi|_n \in P(A)$ and Claim 1 is proved.

To prove Claim 2, fix $\sigma \in P(A)$. We will in fact prove the stronger property that $\sum_{\pi:\pi|_n=\sigma} D_\pi = 1$. Consider a permutation $\pi \in S_{2n}$ with $\pi|_n = \sigma$. By the argument above, each element $h_{i,\sigma i}$ corresponds to two factors in D_π: either $b_{i,\sigma i}$ and $\overline{b}_{i,\sigma i}$ or $c_{i,\sigma i}$ and $-\overline{c}_{i,\sigma i}$. Thus the set $\{\pi : \pi|_n = \sigma\}$ is in 1-1 correspondence with the subsets of $[n]$, where the subset specifies those i which contribute factors $b_{i,\sigma i}$. Recall that for each i, either $b_{i,\sigma i} \in \mathbb{C}_4$ and $c_{i,\sigma_i} = 0$ or vice versa. Hence $D_\pi = 0$ for all but one of these permutations π, namely the permutation $\hat{\pi}$ corresponding to the subset $N = \{i : b_{i,\sigma i} \neq 0\}$. For this permutation, we have $D_{\hat{\pi}} = \mathrm{sgn}(\hat{\pi})(-1)^{n-|N|}$. And an easy induction on $n - |N|$ establishes that $\mathrm{sgn}(\hat{\pi}) = (-1)^{n-|N|}$, from which Claim 2 follows.

This concludes the proof of the first part of Theorem 4.2. □

4.3 Analysis of second moment

To prove the second moment claim in Theorem 4.2, we will proceed in similar fashion to the previous section. In particular, the main step once again is to express the second moment of the estimator for a general matrix A in terms of that for the 2×2 all-1's matrix A_1:

THEOREM 4.3. *Let $E_1 = \mathrm{E}[Y_{A_1}^2]$. Then for any $n \times n$ $(0,1)$ matrix A, we have*

$$\mathrm{E}[Y_A^2] = \sum_{G \in \mathcal{G}(A)} (E_1 - 2)^{c(G)}.$$

The proof of Theorem 4.3 is deferred to the full version of the paper. It is similar in overall structure to our second-moment analysis for general CL_m in section 3, but both simpler because of the fixed dimension $2^{m-1} = 4$ and more complex because of the Dieudonné determinant. This latter complication is handled with similar technology to the expectation analysis in the proof we have just given.

In light of Theorem 4.3, it remains only to compute E_1. Let $H = (h_{ij})$ be the random quaternion matrix computed by the algorithm when run on A_1. Following the progress of the algorithm Gauss shows that

$$Y_{A_1} = |\mathrm{Gauss}(H)|^2 = |h_{11}h_{22} - h_{11}h_{21}\overline{h}_{11}h_{12}|^2.$$

Thus Y_{A_1} has the same distribution as $|h - g|^2$, where h, g are chosen independently and u.a.r. from $\mathbb{H}_8$. An easy hand calculation now shows that Y_{A_1} takes the values 0 and 4 each with probability $\frac{1}{8}$, and the value 2 with probability $\frac{3}{4}$. Thus $E_1 = 5$, so from Theorem 4.3 we get $\mathrm{E}[Y_A^2] = \sum_{G \in \mathcal{G}(A)} 3^{c(G)}$ for an arbitrary A. Hence the critical ratio is bounded above by

$$\frac{\mathrm{E}[Y_A^2]}{\mathrm{E}[Y_A]^2} = \frac{\sum_{G \in \mathcal{G}(A)} 3^{c(G)}}{\sum_{G \in \mathcal{G}(A)} 2^{c(G)}} \leq \max_{G \in \mathcal{G}(A)} \left(\frac{3}{2}\right)^{c(G)} \leq \left(\frac{3}{2}\right)^{n/2}.$$

This completes the proof of the second claim in Theorem 4.2, and the analysis of our modified quaternion estimator. □

4.4 Beyond the quaternions

The question of whether the performance of our higher-dimensional Clifford algebra estimators X_A can also be achieved in polynomial time remains an intriguing open problem. Of course, a positive resolution would in light of Theorem A imply a fully-polynomial randomized approximation scheme for the permanent completely different from that of [8]. We have developed an efficient version of the estimator in the next algebra, CL_4, which according to large-scale experiments has the same behavior as the corresponding X_A. This would actually overcome another major obstacle, as CL_4 is the first algebra that is not a division algebra (i.e., not all elements have inverses).[§] However, so far we have not been able to go beyond this. We note that CL_m for any m has a representation as $k \times k$ matrices over $\mathbb{R}$, $\mathbb{C}$ or $\mathbb{H}$ (or direct sums of such matrices), so one can always define an analog of the "reduced norm" we used for the quaternions. It is also easy to see that the resulting estimator is unbiased, but experiments indicate that the variance is much larger than that of the corresponding X_A.

ACKNOWLEDGEMENTS

We would like to thank Amir Jafari, Laci Lovász, Mike Luby, Bernd Sturmfels and especially T-Y. Lam for helpful discussions.

REFERENCES

[1] E. ARTIN, *Geometric Algebra*, Wiley Interscience, New York, 1988.

[2] A. BARVINOK, Polynomial time algorithms to approximate permanents and mixed discriminants within a simply exponential factor, *Random Structures and Algorithms* **14** (1999), 29–61.

[3] A. BARVINOK, "New permanent estimators via noncommutative determinants," Preprint, July 2000, available from
www.math.lsa.umich.edu/~barvinok/papers.html.

[4] A.Z. BRODER, 'How hard is it to marry at random? (On the approximation of the permanent)," *Proceedings of the 18th Annual ACM Symposium on Theory of Computing* (STOC), ACM Press, 1986, 50–58. Erratum in *Proceedings of the 20th Annual ACM Symposium on Theory of Computing*, 1988, p. 551.

[§] Note that the quaternions are known to have the highest dimension among all division algebras over the reals.

[5] C. Godsil and I. Gutman, “On the matching polynomial of a graph,” *Algebraic Methods in Graph Theory*, 1981, pp. 241–249.

[6] A. Frieze and M. Jerrum, “An analysis of a Monte Carlo algorithm for approximating the permanent,” *Combinatorica* **15** (1995), pp. 67–83.

[7] M. Jerrum and A. Sinclair, “Approximating the permanent,” *SIAM Journal on Computing* **18** (1989), 1149–1178.

[8] M. Jerrum, A. Sinclair and E. Vigoda, “A polynomial-time approximation algorithm for the permanent of a matrix with non-negative entries,” *Proceedings of the 33rd ACM Symposium on Theory of Computing*, 2001, pp. 712–721.

[9] M. Jerrum and U. Vazirani, “A mildly exponential approximation algorithm for the permanent,” *Algorithmica* **16** (1996), 392–401.

[10] N. Karmarkar, R. Karp, R. Lipton, L. Lovász and M. Luby, “A Monte-Carlo algorithm for estimating the permanent,” *SIAM Journal on Computing* **22** (1993), pp. 284–293.

[11] C. Kenyon, D. Randall and A. Sinclair, “Approximating the number of dimer coverings of a lattice,” *Journal of Statistical Physics* **83** (1996), pp. 637–659.

[12] T.-Y. Lam, *The algebraic theory of quadratic forms*, Benjamin/Addison-Wesley, Reading, MA, 1973. (Reprinted with revisions, 1980.)

[13] N. Linial, A. Samorodnitsky and A. Wigderson, “A deterministic strongly polynomial algorithm for matrix scaling and approximate permanents,” *Combinatorica* **20** (2000), 545–568.

[14] H. Minc, *Permanents*, Encyclopedia of Mathematics and its Applications **6** Addison-Wesley Publishing Company, 1982.

[15] N. Nisan, “Lower bounds for non-commutative computation,” *Proceedings of the 23rd Annual ACM Symposium on Theory of Computing*, 1991, pp. 410–418.

[16] L.E. Rasmussen, “Approximating the permanent: a simple approach,” *Random Structures and Algorithms* **5** (1994), pp. 349–361.

[17] L.E. Rasmussen, “On approximating the permanent and other #P-complete problems,” PhD Thesis, Computer Science Division, UC Berkeley, 1998.

[18] L.G. Valiant, “The complexity of computing the permanent,” *Theoretical Computer Science* **8** (1979), 189–201.

[19] B.L. van der Waerden, *Algebra* Vol. 2, Frederick Ungar Publishing Co., New York, 1970.

APPENDIX

A.2. THE CLIFFORD ALGEBRA CL_4

The Clifford algebra CL_4 has eight basis elements, namely $\{1, u_{12}, u_{23}, u_{13}, u_{1234}, u_{34}, u_{14}, u_{24}\}$. The only self-conjugate basis elements are 1 and u_{1234}. Note that these are also the only two elements that commute with all others. The complete multiplication table is as follows:

$$\begin{bmatrix} 1 & u_{12} & u_{23} & u_{13} & u_{1234} & u_{34} & u_{14} & u_{24} \\ u_{12} & -1 & u_{13} & -u_{23} & -u_{34} & u_{1234} & -u_{24} & u_{14} \\ u_{23} & -u_{13} & -1 & u_{12} & -u_{14} & u_{24} & u_{1234} & -u_{34} \\ u_{13} & u_{23} & -u_{12} & -1 & u_{24} & u_{14} & -u_{34} & -u_{1234} \\ u_{1234} & -u_{34} & -u_{14} & u_{24} & 1 & -u_{12} & -u_{23} & u_{13} \\ u_{34} & u_{1234} & -u_{24} & -u_{14} & -u_{12} & -1 & u_{13} & u_{23} \\ u_{14} & u_{24} & u_{1234} & u_{34} & -u_{23} & -u_{13} & -1 & -u_{12} \\ u_{24} & -u_{14} & u_{34} & -u_{1234} & u_{13} & -u_{23} & u_{12} & -1 \end{bmatrix}$$

If $h = c_1 + c_2u_{12} + c_3u_{23} \underline{+} c_4u_{13} + c_5u_{1234} + c_6u_{34} + c_7u_{14} + c_8u_{24}$ then its conjugate $\overline{h}$ is defined as

$$\overline{h} = c_1 - c_2u_{12} - c_3u_{23} - c_4u_{13} + c_5u_{1234} - c_6u_{34} - c_7u_{14} - c_8u_{24}.$$

Note that

$$h\overline{h} = \sum\nolimits_{i=1}^{8} c_i^2 + 2(c_1c_5 - c_2c_6 - c_3c_7 + c_4c_8)u_{1234},$$

which is not real. The norm-square is defined by $|h|^2 = \sum_{i=1}^{8} c_i^2$.

A.3. SOME PROOFS FROM SECTION 3

In this section we fill in the proof details for Lemmas 3.3 and 3.5 that were omitted from the main text.

Proof of Lemma 3.3: Before proceeding with the case analysis of equation (2) outlined in the main text, we require some easy facts about self-conjugate subgroups.

Lemma A.3.1. *Let G be a subgroup of G_m, and let a be any element of G_m. Let H be the subset of G that commutes with a, i.e., $b \in H$ if and only if $ab = ba$. Then H is a subgroup of G, and furthermore, either $H = G$ or $|H| = |G|/2$.*

Proof: That H is a subgroup of G is immediate. Suppose that $H \neq G$, and let $G - H = \{b_1, \dots b_r\}$ be those elements of G that do not commute with a. Now notice any product b_ib_j does belong to H. Thus the elements b_1b_i are all distinct and belong to H, so H is at least as large as $G - H$. □

Lemma A.3.2. *In the situation of the previous lemma, with $|H| = |G|/2$, let g be an element of G but not H. Then $gH = G - H$.*

Lemma A.3.3. **[Expansion]** *Let G be a self-conjugate subgroup of G_m, and let c be a self-conjugate element of G_m that commutes with G but is not in $G \cup -G$. Then $G \cup cG$ is also a self-conjugate subgroup of G_m and has twice the size of G.*

Lemma A.3.4. **[Conjugation]** *Let G be an arbitrary subgroup of G_m. For any $a \in G_m$, $a(\sum_{\alpha \in G} u_\alpha)\overline{a}$ can be written as $\sum_{\alpha \in G'} u_\alpha$, where $G' = aG\overline{a}$ is a conjugate subgroup of G. Hence if G is self-conjugate, then so is G'.*

We now proceed with the case analysis from the main part of the paper.

Case 1: c commutes with all of G.

Here (2) becomes $2^k a(2\sum_{\alpha \in G} u_\alpha + (c+\overline{c})\sum_{\alpha \in G} u_\alpha)\overline{a}$. We now analyze the three subcases outlined in the main text:

case 1a: $c = \overline{c}$, $c \in G \cup -G$. Observe that (2) becomes $2^k a(2\sum_{\alpha \in G} u_\alpha + 2c\sum_{\alpha \in G} u_\alpha)\overline{a}$. Now if $c \in G$ then $cG = G$, so $c\sum_{\alpha\in G} u_\alpha = \sum_{\alpha\in G} u_\alpha$ and we get $2^{k+2}a(\sum_{\alpha\in G} u_\alpha)\overline{a}$. We then apply the conjugation lemma. Similarly, if $c \in -G$, then $cG = -G$ and we end up with 0.

CASE 1B: $c = \overline{c}$, $c \notin G \cup -G$. Again, (2) becomes $2^k a(2\sum_{\alpha\in G} u_\alpha + 2c\sum_{\alpha\in G} u_\alpha)\overline{a}$. Since c is self-conjugate and commutes with G, from the expansion lemma we can write this as $2^{k+1}a(\sum_{\alpha\in G'} u_\alpha)\overline{a}$ where $G' = G \cup cG$ is also self-conjugate. We then apply the conjugation lemma.

CASE 1C: $c = -\overline{c}$. Here (2) becomes $2^{k+1}a(\sum_{\alpha\in G} u_\alpha)\overline{a}$, and the conjugation lemma finishes this case.

CASE 2: c does not commute with all of G.

By Lemma A.3.1, the set of elements that commute with c form a subgroup $H \subset G$ with $|H| = |G|/2$. In this case, the first two terms in (2) become $\sum_{\alpha\in G} u_\alpha + c(\sum_{\alpha\in G} u_\alpha)\overline{c} = 2\sum_{\alpha\in H} u_\alpha$. Analyzing the last two terms will require two subcases:

CASE 2A: $c = \overline{c}$. Here we get $c(\sum_{\alpha\in G} u_\alpha) + (\sum_{\alpha\in G} u_\alpha)\overline{c} = 2c\sum_{\alpha\in H} u_\alpha$, so upon combining with the first two terms (2) becomes $2^{k+1}a(\sum_{\alpha\in H} u_\alpha + c\sum_{\alpha\in H} u_\alpha)\overline{a}$. H must be self-conjugate since it is a subgroup of G. c is self-conjugate, commutes with H, and does not belong to $H \cup -H$. Thus the expansion lemma allows us to rewrite $\sum_{\alpha\in H} u_\alpha + c\sum_{\alpha\in H} u_\alpha$ as $\sum_{\alpha\in G'} u_\alpha$ where $G' = H \cup cH$ is self-conjugate and abelian. We then apply the conjugation lemma and we are done.

CASE 2B: $c = -\overline{c}$. Here $c(\sum_{\alpha\in G} u_\alpha) + (\sum_{\alpha\in G} u_\alpha)\overline{c} = 2c\sum_{\alpha\in G-H} u_\alpha$, so upon combining with the first two terms (2) becomes $2^{k+1}a(\sum_{\alpha\in H} u_\alpha + c\sum_{\alpha\in G-H} u_\alpha)\overline{a}$. Recall from Lemma A.3.2 that $G - H = gH$ for some $g \in G$ not in H. Thus we can rewrite $\sum_{\alpha\in H} u_\alpha + c\sum_{\alpha\in G-H} u_\alpha$ as $\sum_{\alpha\in H} u_\alpha + cg\sum_{\alpha\in H} u_\alpha$. Since cg is self-conjugate ($(cg)^2 = cgcg = -ccgg = -c^2g^2 = -(-1)(1) = 1$) and commutes with H (both c and g commute with H), we can again apply the expansion lemma followed by the conjugation lemma.

This completes the proof of Lemma 3.3. □

Proof of Lemma 3.5: Here we prove the stronger version of Lemma 3.5 stated in the main text, namely that $p = \frac{1}{2^{2q+1}}$. We first require two lemmas:

LEMMA A.3.5. *Let $m \not\equiv 0 \bmod 4$. Suppose a subgroup H of size 2^{k-1} expands to a subgroup G of size 2^k. Then $|Z(H)| = 2|Z(G)|$, where $Z(H)$ and $Z(G)$ are the centralizers of H and G in G_m.*

Proof: Recall that the *centralizer* of a subgroup G in G_m is the set of elements of G_m that commute with all of G.

The proof appeals to a linear algebra description of G_m. Consider $\mathbb{F}_2^m$, the m-dimensional vector space over $\mathbb{F}_2$. We identify a basis element a of G_m with the vector v such that v_i is 1 if and only if i appears among the subscripts of a. Note that given two basis elements a and b their product (up to sign) is described by the sum of their corresponding vectors $v + w$. Further, if we define the dot product in the usual way, then a and b commute if and only if $v \cdot w = 0$.

A self-conjugate subgroup K of size 2^j can then be represented (up to sign) as a j-dimensional subspace W_K of $\mathbb{F}_2^m$. K is generated by exactly j elements, and a basis for the subspace W_K consists of the vectors corresponding to these generators. These vectors are distinct because of our restriction that if $a \in K$ then $a \notin -K$. The vectors must be linearly independent since a linear dependence would imply that one generator is a product of the others up to sign, which is impossible.

We define the subspace V_K to be the subspace spanned by W_K and $\overline{1}$, the all-ones vector. The basis elements that commute with K are then exactly those represented by the orthogonal subspace $V_K^\perp$. Any $v \in V_K^\perp$ must have an even number of subscripts since $v \cdot \overline{1} = 0$ and must commute with K since $v \cdot w = 0$ for all $w \in W_K$. It is a well-known fact that $\dim V_K + \dim V_K^\perp = \dim \mathbb{F}_2^m = m$. There are thus $2^{\dim V_K^\perp}$ unsigned basis elements in $Z(K)$ so $|Z(K)| = 2^{\dim V_K^\perp + 1}$

Finally, observe that if H expands to G, then $\dim V_H = \dim V_G - 1$. This is true since G requires one more generator than H, and further, none of these generators can be $\overline{1}$ since $m \not\equiv 0 \bmod 4$. Thus $|Z(H)| = 2|Z(G)|$. □

LEMMA A.3.6. *When H expands to G as above, exactly half of the elements in $Z(H) - Z(G)$ are self-conjugate.*

Proof: Recall from case 1B of the proof of Lemma 3.3 that G can be described as $a(H \cup cH)\overline{a}$ for some self-conjugate c that commutes with H. Since conjugation by a changes only the signs of $H \cup cH$ we see that $Z(G) = Z(H \cup cH)$.

Note that $H \subseteq Z(H)$, and consider the set of cosets of H in $Z(H)$. Note that the elements of a coset are either all self-conjugate or all not self-conjugate ($(dh)^2 = dhdh = ddhh = d^2$). Also, the elements of a coset either all commute with c (and hence with G) or all do not commute with c ($dhc = cdh \Leftrightarrow dch = cdh \Leftrightarrow dc = cd$).

Now consider a coset dH in $Z(H) - Z(G)$ that is not self-conjugate. Then cdH is also in $Z(H) - Z(G)$ but is self-conjugate. Similarly if dH in $Z(H) - Z(G)$ is self-conjugate, then cdH is not. Thus we have a bijection between self-conjugate and non-self-conjugate cosets in $Z(H) - Z(G)$, proving our result. □

We now use the above two lemmas to show that when $m = 4q + 2$, the largest self-conjugate subgroup has size exactly 2^{2q}.

When $m \equiv 2 \bmod 4$, the self-conjugate elements comprise exactly half of the 2^{m-1} unsigned basis elements. (To see this, recall that the unsigned basis elements correspond to even cardinality subsets of $\{1, \ldots, m\}$; the self-conjugate elements correspond to those subsets of cardinality $4k + 2$.) Thus when the subgroup has size $2^0 = 1$, its centralizer contains 2^{m-2} self-conjugate unsigned basis elements. At each expansion step, the subgroup size doubles and the number of self-conjuate unsigned basis elements in the centralizer halves until the two are the same; this occurs when the subgroup has size 2^t, where $t = \frac{m-2}{2} = 2q$.

Thus $p = \max_G \frac{2|G|}{|G_m|}$ is exactly $\frac{1}{2^{2q+1}}$, as desired. □