

HOLOGRAPHIC ALGORITHMS WITHOUT MATCHGATES

J.M. LANDSBERG, JASON MORTON AND SERGUEI NORINE

ABSTRACT. The theory of *holographic algorithms*, which are polynomial time algorithms for certain combinatorial counting problems, surprised the complexity community by showing certain problems, very similar to $\#\mathbf{P}$ complete problems, were in fact in the class \mathbf{P} . In particular, the theory produces algebraic tests for a problem to be in \mathbf{P} . In this article we describe the geometric basis of these algorithms by (i) replacing the construction of graph fragments in the procedure by the direct construction of a skew symmetric matrix, and (ii) replacing the computation of weighted perfect matchings of an auxiliary graph by computing the Pfaffian of the directly constructed skew-symmetric matrix. This procedure indicates a more geometric approach to complexity classes. It also leads to more general constructions where one replaces the “Grassmann-Plücker identities” which test for admissibility by other algebraic tests. Natural problems treatable by these methods have been previously considered in a different context, and we present one such example.

1. INTRODUCTION

1.1. **History.** In [19, 20, 21, 23, 24, 25] L. Valiant introduced *matchgates* and *holographic algorithms*, in order to prove the existence of polynomial time algorithms for counting and sum-of-products problems that naïvely appear to have exponential complexity. Such algorithms have been studied in depth and further developed by J. Cai et al. [1, 2, 3, 4, 5, 6, 7, 8].

The determinant of an $n \times n$ matrix, expressed as a sum of terms indexed by the symmetric group S_n , has $n!$ terms, yet it can be computed in polynomial time. Holographic algorithms are a particular strategy for finding algorithms which provide efficient computation of such exponential sums of products by essentially expressing the problem as a determinant. Some exploit a “Hadamard-type” change of basis inspired by quantum computation and this “superposition,” together with the cancellation effect present in the determinant, is the source of the name. Holographic algorithms have also been called “accidental algorithms” because they sometimes provide a polynomial time algorithm for certain parameter values in a parameterized family of problems, and these problems are \mathbf{NP} -hard (e.g. because they are $\oplus P$ complete) or $\#\mathbf{P}$ -complete for other parameter values.

For the reader not familiar with complexity theory, a typical problem considered in complexity theory is as follows: let x_1, \dots, x_n be a collection of Boolean variables (i.e., variables assuming the values “true” or “false”) and let c_1, \dots, c_p be a collection of clauses (e.g., $x_1 \wedge (\neg x_3)$). Then one wants to count the number of satisfying assignments. A geometer might wish to view the x_j as taking values in the field with two elements \mathbb{F}_2 , the c_s as equations on \mathbb{F}_2^n , and the problem is to count the number of points of the variety determined by the c_s . Such a problem is called $\#\text{SAT}$ in the complexity literature. (The $\#$ refers to counting, if it is left off, the problem is just to determine if there exists any satisfying assignment.)

Often restrictions are placed on the types of clauses allowed. Typical restrictions (with their names) are

- (1) 3SAT: Each clause involves only three variables.

- (2) CNF: conjunctive normal form, i.e., a collection of “or” clauses joined together with a series of “ands” e.g., $(x_1 \vee x_3) \wedge (x_2 \vee x_5 \vee x_4) \wedge (x_7)$.
- (3) Mon: restriction to monotone clauses - variables appearing in a clause either all appear positively or all negatively.
- (4) Rtw: read twice - each variable appears in just two clauses.
- (5) NAE: all clauses are “not all equal”- variables appearing in the same clause are not all allowed to have the same assignments.
- (6) Pl: planar: the variable-clause inclusion graph (see below) constructed from each instance is planar

Problems are named after these restrictions, for example $\#$ Pl-Rtw-Mon-3CNF (defined in [22]) is to count the number of satisfying assignments when clauses are subject to the restrictions Pl-Rtw-Mon-3CNF. For this problem, even counting the number of solutions mod 2 is $\oplus\mathbf{P}$ complete (this problem is called $\#_2$ Pl-Rtw-Mon-3CNF) and therefore \mathbf{NP} -hard by Valiant-Vazirani reduction [26]. Nevertheless, a holographic algorithm shows that counting the number of solutions to $\#$ Pl-Rtw-Mon-3CNF mod 7 (called $\#_7$ Pl-Rtw-Mon-3CNF) can be accomplished in polynomial time [22].

Remark 1.1. Note that all the conditions above except the last are “local”. The global nature of the last condition turns out to be crucial for implementation of holographic algorithms.

The perspective provided by holographic algorithms relates the computational complexity of certain key problems with the solvability of systems of polynomial equations [22]. A popular account describing holographic algorithms can be found in [11].

Holographic algorithms work as follows: suppose the problem \mathcal{P} is SAT or some restricted version of SAT. Say an instance P of \mathcal{P} has variables x_1, \dots, x_m subject to clauses c_1, \dots, c_p .

- (1) Associate to P a bipartite graph $\Gamma_P = (V, U, E)$, with vertex sets $V = \{x_1, \dots, x_m\}$ and $U = \{c_1, \dots, c_p\}$. Draw an edge $(i, s) \in E$ iff x_i appears in c_s .
- (2) Associate (two dimensional) vector spaces A_e , equipped with bases, to each edge, “local” tensors to each vertex, and a “global” tensor to each vertex set, such that the two global tensors lie in dual vector spaces. Write $G \in \otimes_{e \in E} A_e \sim \mathbb{C}^{2^{|E|}}$, $R \in \otimes_{e \in E} A_e^*$ for the global tensors. The construction is such that the pairing $\langle G, R \rangle \in \mathbb{C}$ counts the number of satisfying assignments (so the pairing actually takes values in the nonnegative integers). This step is explained in §2.
- (3) Perform a (cheap) basis change in each A_e and the dual basis change in each A_e^* such that local *matchgate identities* are satisfied. The matchgate identities are a collection of polynomial equations in the coordinates of the clauses and variables. This step is explained in §3.
- (4) Replace the vertices of Γ_P with weighted graph fragments to obtain a new weighted graph $\Gamma_{\Omega(P)}$ such that the weighted sum of perfect matchings of $\Gamma_{\Omega(P)}$ equals $\langle G, R \rangle$. The weighted graph fragments are called *matchgates*.
- (5) Assuming $\Gamma_{\Omega(P)}$ is planar (or more generally *Pfaffian*), compute the weighted sum of perfect matchings of $\Gamma_{\Omega(P)}$ quickly using the FKT algorithm [14, 18]. The FKT algorithm for weighted planar graphs is a way to alter signs of weights to convert the Hafnian (which counts the number of perfect matchings of a weighted graph), to a Pfaffian, which is equivalent in complexity to the determinant.

In the early matchgates literature the construction of the matchgate graph fragments was viewed as an “art”, but since then, a library of such constructions has been made by Cai et. al. at least for vertices without too many edges, so it is also essentially algorithmic to apply. The matchgate identities are equalities among sums of products of sub-Pfaffians. We discuss them more below.

1.2. Holographic algorithms without matchgates. Our approach, described in §4, replaces steps (4) and (5) with the computation of the Pfaffian of a natural $|E| \times |E|$ matrix associated to the original graph Γ_P . We present an algorithmic construction of a small matrix for each variable and clause, which is the analog of the matchgate construction.

We replace FKT with an edge ordering defined by a plane curve as described in Section 5. Both FKT and our edge ordering have the effect of “making the signs work out right”.

The Valiant-Cai formulation of holographic algorithms can be summarized as

$$(1) \quad \#\text{satisfying assignments of } P = \langle G, R \rangle = \text{weighted sum of perfect matchings of } \Gamma_{\Omega(P)}.$$

We associate constants $\alpha = \alpha_G, \beta = \beta_R$ (depending only on the number of each type of vertex) and $|E| \times |E|$ -skew symmetric matrices $\tilde{z} = \tilde{z}_G, y = y_R$ directly to G, R , *without the construction of matchgates*, to obtain the equality:

$$(2) \quad \#\text{satisfying assignments of } P = \langle G, R \rangle = \alpha\beta\text{Pfaff}(\tilde{z} + y);$$

see Examples 4.8 and 4.9. The constants and matrices are essentially just components of the vectors G, R . The algorithm complexity is dominated by evaluating the Pfaffian of the $|E| \times |E|$ skew-symmetric matrix.

The key to our approach is that a vector satisfies the matchgate identities iff it is a vector of sub-Pfaffians of some skew-symmetric matrix, and that the pairing of two such vectors can be reduced to calculating a Pfaffian of a new matrix constructed from the original two. A similar phenomenon holds in greater generality discussed in Appendix 7. A simple example is the set of vectors of sub-minors of an arbitrary rectangular matrix. We describe an example of such an implementation in Section 6.

The starting point of our investigations was the observation that the matchgate identities come from classical geometric objects called *spinors*. The results in this article do not require any reference to spinors to either state or prove, and for the convenience of the reader not familiar with them we have eliminated all mention of them except for this paragraph and an Appendix (§7) included for the interested reader.

The purpose of this paper is to facilitate an understanding of the geometric basis of holographic algorithms. It does not address such issues as implementation, the construction of explicit new algorithms, or dependence on the ground field. The first two topics are addressed in [15]. The third is beyond our expertise, but since the circulation of this paper in preprint form, A. Snowden has begun to work on this issue.

2. COUNTING PROBLEMS AS TENSOR CONTRACTIONS

For brevity we continue to restrict to problems counting the number of satisfying assignments of Boolean variables x_i subject to clauses c_s (such as #PI-Mon-NAE-SAT defined above). Following e.g., [2] express an instance P of such a problem in terms of a tensor contraction diagrammed by a planar bipartite graph $\Gamma_P = (V, U, E)$ as above (see Figure 1), together with the data of tensors $G_i = G_{x_i}$ and $R_s = R_{c_s}$ attached at each vertex $x_i \in V$ and $c_s \in U$. G_i will record that x_i is 0 or 1 and R_s will record that the clause c_s is satisfied. Let $n = |E|$ be the number of edges in Γ_P .

For each edge $e = (i, s) \in E$ define a 2-dimensional \mathbb{C} vector space A_e with basis $a_{e|0}, a_{e|1}$. Say x_i has degree d_i and is joined to $c_{j_1}, \dots, c_{j_{d_i}}$. Let E_i denote the set of edges incident to x_i and associate to each x_i the tensor

$$(3) \quad G_i := a_{i,s_{j_1}|0} \otimes \cdots \otimes a_{i,s_{j_{d_i}}|0} + a_{i,s_{j_1}|1} \otimes \cdots \otimes a_{i,s_{j_{d_i}}|1} \in A_i := A_{i,s_{j_1}} \otimes \cdots \otimes A_{i,s_{j_{d_i}}} \\ = \bigotimes_{e \in E_i} a_{e|0} + \bigotimes_{e \in E_i} a_{e|1}$$

The tensor G_i represents that either x_i is true (all 1's) or false (all 0's), and presents the same value to each clause. It is called a *generator* in the matchgates literature and is denoted by the co-ordinate vector $(1, 0, \dots, 0, 1)$ corresponding to a lexicographic basis of $A_i = \otimes_{e \in E_i} A_e$. This vector is called its *signature*. We use notation emphasizing the tensor product structure of the vector space $A_i = \mathbb{C}^{2^{d_i}}$, and will use the word signature to refer to the tensor expression of G_i .

Next define a tensor associated to each clause c_s representing that c_s is satisfied. Let A_e^* be the dual space to A_e with dual basis $\alpha_{e|0}, \alpha_{e|1}$. Let E_s denote the set of edges incident to c_s . For example, if c_s has degree d_s and is “not all equal” (NAE), then the corresponding tensor (called a *recognizer*) associated to it is

$$(4) \quad R_s := \sum_{(\epsilon_1, \dots, \epsilon_{d_s}) \neq (0, \dots, 0), (1, \dots, 1)} \alpha_{i, s_1 | \epsilon_1} \otimes \dots \otimes \alpha_{i, s_{d_s} | \epsilon_{d_s}} = \sum_{(\epsilon_1, \dots, \epsilon_{d_s}) \neq (0, \dots, 0), (1, \dots, 1)} \bigotimes_{e \in E_s} \alpha_{e | \epsilon_e}$$

Remark 2.1. Our restriction to $\#$ -NAE-3SAT is to facilitate presentation. It is already $\#$ P hard.

Now consider $G := \otimes_i G_i$ and $R := \otimes_s R_s$ respectively elements of the vector spaces $A := \otimes_e A_e$ and $A^* := \otimes_e A_e^*$. Then the number of satisfying assignments to P is $\langle G, R \rangle$ where $\langle \cdot, \cdot \rangle : A \times A^* \rightarrow \mathbb{C}$ is the pairing of dual vector spaces. At this point we have merely exchanged our original counting problem for the computation of a pairing in vector spaces of dimension $2^{|E|}$.

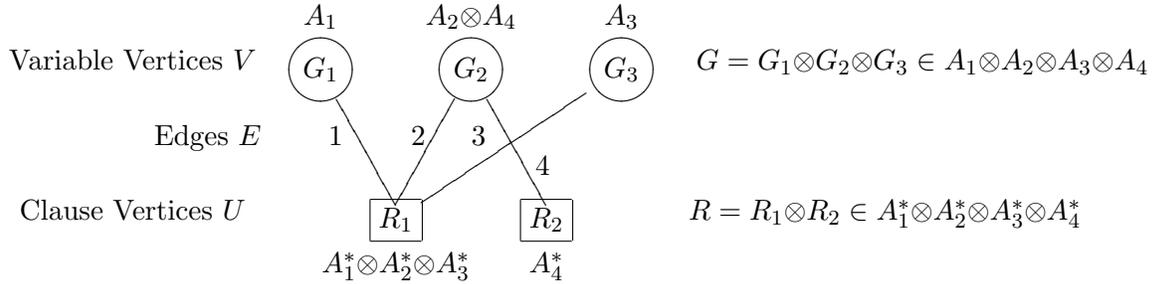


FIGURE 1. A bipartite graph Γ diagrams a tensor contraction, representing an exponential sum of products such as counting the satisfying assignments of a satisfiability problem. Boxes denote clauses, circles denote variables. Each clause or variable corresponds to a tensor lying in the indicated vector space; e.g. $R_1 \in A_1^* \otimes A_2^* \otimes A_3^*$. Instead of replacing each vertex with a matchgate, our construction defines an $n \times n$ matrix, where n is the number of edges in the problem graph. The Pfaffian of this matrix, times a constant depending on the number of each type of variable and clause, is the number of satisfying assignments.

3. LOCAL CONDITIONS AND CHANGE OF BASIS

In order to be able to construct the matchgates corresponding to the x_i, c_s , there are local conditions that need to be satisfied; the algebraic equations placed on the G_i, R_s are called the *Grassmann-Plücker identities* (or *Matchgate Identities* in this context). See, e.g., Theorem 7.2 of [4] for an explicit expression of the equations, which date back at least to Chevalley in the 1950's [9]. From our perspective, these identities ensure that a tensor T representing a variable or clause can be written as a vector of sub-Pfaffians of some matrix. From the matchgates point of view, these equations are necessary and sufficient conditions for the existence of graph fragments that can replace the vertices of Γ_P to form a new weighted graph $\Gamma_{\Omega(P)}$ such that the weighted perfect matching polynomial of $\Gamma_{\Omega(P)}$ equals $\langle G, R \rangle$.

Expressed in the basis most natural for a problem, a clause or variable tensor may fail to satisfy the Grassmann-Plücker identities. However it may do so under a change of basis; e.g. in Example 3.1 below, we replace the basis (True, False) with (True+False, False−True). Such a change of basis will not change the value of the pairing $A \times A^* \rightarrow \mathbb{C}$ as long as we make the corresponding dual change of basis in the dual vector space—but of course this may cause the tensors in the dual space to fail to satisfy the identities. Thus one needs a change of basis that works for both generators and recognizers. In this article, as in almost all existing applications of the theory, we only consider changes of bases in the individual A_e 's, and we will perform the exact same change of basis in each such, although neither restriction is *a priori* necessary for the theory.

Remark 3.1. While it may seem incredible that such changes of basis might exist for a problem, in fact for many NP-hard problems there exist local changes of basis that work for the G_i and for the R_s . Moreover the local changes for the G_i can be glued together to make a global change that works for the tensor G , and the local changes for the R_s can be glued together to make a global change that works for the tensor R . The “only” problem that occurs is a compatibility of signs, which needs a global condition such as planarity to be overcome.

3.1. Example. In #Mon-3-NAE-SAT the generator tensor G_i corresponding to a variable vertex x_i is (3). The recognizer tensor corresponding to a NAE clause R_s is (4) and in our case we will have $d_s = 3$ for all s .

Let T_0 be the basis change, the same in each A_e , sending $a_{e|0} \mapsto a_{e|0} + a_{e|1}$ and $a_{e|1} \mapsto a_{e|0} - a_{e|1}$ which induces the basis change $\alpha_{e|0} \mapsto \frac{1}{2}(\alpha_{e|0} + \alpha_{e|1})$ and $\alpha_{e|1} \mapsto \frac{1}{2}(\alpha_{e|0} - \alpha_{e|1})$ in A_e^* . This basis is denoted **b2** in [25]. Applying T_0 , we obtain

$$T_0(a_{i,s_{i_1}|0} \otimes \cdots \otimes a_{i,s_{i_{d_i}}|0} + a_{i,s_{i_1}|1} \otimes \cdots \otimes a_{i,s_{i_{d_i}}|1}) = 2 \sum_{\{(\epsilon_1, \dots, \epsilon_{d_i}) \mid \sum \epsilon_\ell = 0 \pmod{2}\}} a_{i,s_{i_1}|\epsilon_1} \otimes \cdots \otimes a_{i,s_{i_{d_i}}|\epsilon_{d_i}}.$$

In the matchgates literature this tensor is denoted by the vector $(2, 0, 2, 0, \dots, 2, 0, 2)$ (assuming the number of incident edges is even). The action on recognizer tensors is

$$\begin{aligned} T_0 \left(\sum_{(\epsilon_1, \epsilon_2, \epsilon_3) \neq (0,0,0), (1,1,1)} \alpha_{i,s_1|\epsilon_1} \otimes \alpha_{i,s_2|\epsilon_2} \otimes \alpha_{i,s_3|\epsilon_3} \right) \\ = 6\alpha_{i,s_1|0} \otimes \alpha_{i,s_2|0} \otimes \alpha_{i,s_3|0} - 2(\alpha_{i,s_1|0} \otimes \alpha_{i,s_2|1} \otimes \alpha_{i,s_3|1} + \alpha_{i,s_1|1} \otimes \alpha_{i,s_2|0} \otimes \alpha_{i,s_3|1} + \alpha_{i,s_1|1} \otimes \alpha_{i,s_2|1} \otimes \alpha_{i,s_3|0}) \end{aligned}$$

or, denoted by its co-ordinate vector, $(6, 0, 0, 0, -2, -2, -2, 0)$.

4. HOLOGRAPHIC ALGORITHMS WITHOUT MATCHGATES

4.1. Detour: Elementary identities from linear algebra. Our approach to holographic algorithms essentially boils down to two simple identities from linear algebra. First, if x, y are matrices whose product is a square matrix, then

$$\det(Id + xy)$$

may be expressed as a sum of products of determinants of the square submatrices of x and y .

Second is the even more elementary identity for block diagonal matrices

$$\det \begin{pmatrix} X & 0 \\ 0 & Y \end{pmatrix} = \det(X) \det(Y).$$

More precisely, we use the Pfaffian analogs of these identities. The first is the key to performing the pairing $A \times A^* \rightarrow \mathbb{C}$ mentioned above in polynomial time, the second is the key to passing from the local changes of bases to a global change of basis.

4.2. The complement pairing and representing G and R as vectors of sub-Pfaffians.

Assume we have a problem expressed as in §2 and have constructed tensors G, R such that in some change of basis their component tensors G_i, R_s satisfy the Grassmann-Plücker identities. For the purposes of exposition, we will assume the total number of edges is even. See the discussion in the Appendix §7 for the case of an odd number of edges.

To compute $\langle G, R \rangle$, we will represent G and R as vectors of sub-Pfaffians. For an $n \times n$ skew-symmetric matrix z , the vector of sub-Pfaffians $\text{sPf}(z)$ lies in a vector space of dimension \mathbb{C}^{2^n} , where the coordinates are labeled by subsets $I \subset [n]$, and

$$(\text{sPf}(z))_I = \text{Pfaff}(z_I)$$

where z_I is the submatrix of z including only the rows and columns in the set I . Letting $I^C = [n] \setminus I$, similarly define $\text{sPf}^\vee : \text{Mat}_{n \times n} \rightarrow \mathbb{C}^{2^n}$ by

$$(\text{sPf}^\vee(z))_I = \text{Pfaff}(z_{I^C}).$$

The vector spaces $A = \otimes A_e, A^* = \otimes A_e^*$ come equipped with un-ordered bases induced from the bases of the A_e . These bases do not have a canonical identification with subsets of $(1, \dots, n)$ but do have a convenient choice of identification *after* making a choice of edge ordering. After an ordering \bar{E} of the edges has been chosen we obtain ordered bases of A, A^* . To obtain the convenient choice of identification for A , identify the vector corresponding to $I = (i_1, \dots, i_{2p})$ with the element with 1's in the i_1, \dots, i_{2p} slots and zeros elsewhere, so, e.g. $I = \emptyset$ corresponds to $(0, \dots, 0)$, ..., $I = (1, \dots, n)$ corresponds to $(1, \dots, 1)$. (Recall we are assuming n is even.) Reverse the correspondence for A^* .

For later use, we remark that with these identifications, as long as the first (resp. last) entry of G (resp. R) is non-zero, we may rescale to normalize them to be one. (If say, e.g., the first entry of G is zero but the last is not, and last entry of R is non-zero, we can just reverse the identifications and proceed.) Note that the first and last choices of entries are independent of the edge ordering, but if necessary, to get the first and last entries non-zero, we simply take a less convenient choice of identification. (See §7 for an explanation of this freedom.) As long as this is done consistently it will not produce any problems.

We now explicitly construct the matrices which play the role of matchgates. Suppose G_i is the generator tensor representing a single variable of degree d and that G_i satisfies the matchgate identities. Possibly rescaling G_i by some factor α , we may assume that the coefficient of the lexicographically first term (suppressing the i index in the notation for the rest of this paragraph) $a_{\epsilon_1|0} \otimes \dots \otimes a_{\epsilon_d|0}$ is 1.

Let $1 \leq j < k \leq d$ and consider the $\binom{d}{2}$ generator tensors G_i , with ones ($a_{\epsilon_j|1}$) in the (j, k) -th place and zeros ($a_{\epsilon_r|0}$) elsewhere. Let c_{jk} denote the coefficient of these G_i 's. Let x be the $d \times d$ skew-symmetric matrix with $x_{jk} = c_{jk}$ and $x_{kj} = -c_{jk}$. The matchgate identities, see e.g., [4, Thm. 7.2], imply that for any $(i_1, i_2, \dots, i_d) \in \{0, 1\}^d$, the coefficient of $a_{\epsilon_1|i_1} \otimes \dots \otimes a_{\epsilon_d|i_d}$ in G_i equals the Pfaffian of the principal submatrix of x including the rows where $i_\ell = 1$ and excluding those where $i_\ell = 0$. We write $\text{sPf}(x) = G_i$. The construction is analogous for recognizers. For example, $j = 2, k = 3$ gives the term $c_{23} a_{\epsilon_1|0} \otimes a_{\epsilon_2|1} \otimes a_{\epsilon_3|1} \otimes a_{\epsilon_4|0} \dots \otimes a_{\epsilon_d|0}$.

For the rest of this section, we restrict attention to generator and recognizer tensors at the vertices that can be written $G_i = \text{sPf}(x_i)$ and $R_s = \text{sPf}^\vee(y_s)$ for matrices x_i, y_s , possibly after a change of basis. We also assume for brevity that the G_i and R_s are symmetric, i.e. that $G_i = \text{sPf}(x_i) = \text{sPf}(\pi(x_i))$ for any permutation π on the edges incident on G_i ; this covers many problems of interest. For the more general case when the variables or clauses are not symmetric, and we need to be more careful about defining \bar{E}_G and \bar{E}_R , see §5 and the Appendix §8.

Definition 4.1. Suppose a graph Γ has n edges. An *edge order* is a bijective labelling of the edges of Γ by the integers $1, 2, \dots, n$. Call an edge order such that edges incident on each $x_i \in V$ (resp. $c_s \in U$) have a consecutive block of labels a *generator order* (resp. *recognizer order*) and denote such by \bar{E}_G (resp. \bar{E}_R).

Proposition 4.2. Suppose P is an instance of a problem expressed as a pairing such that \bar{E}_G and \bar{E}_R are respectively generator and recognizer orders to the graph Γ_P . If for all $x_i \in V$ there exists $z_i \in \text{Mat}_{d_i \times d_i}$ such that $\text{sPf}(z_i) = G_i$ under the \bar{E}_G order, and similarly, there exists $y_s \in \text{Mat}_{d_s \times d_s}$ for c_s and R_s with $\text{sPf}^\vee(y_s) = R_s$, then there exists $z, y \in \text{Mat}_{|E| \times |E|}$ such that

$$\begin{aligned} \text{sPf}(z) &= G && \text{with the } \bar{E}_G \text{ order and} \\ \text{sPf}^\vee(y) &= R && \text{with the } \bar{E}_R \text{ order.} \end{aligned}$$

In this situation, z, y are just given by stacking the component matrices z_i, y_s block-diagonally.

Proof. The Pfaffian of a block-diagonal skew-symmetric matrix is the product of the Pfaffians of the diagonal blocks. \square

As the proposition suggests, a difficulty appears when we try to find an order \bar{E} that works for both generators and recognizers.

Definition 4.3. An order \bar{E} is *valid* if there exists skew-symmetric matrices z, y such that $\text{sPf}(z) = G$ and $\text{sPf}^\vee(y) = R$ under the \bar{E} order.

Thus if an order is valid

$$\langle G, R \rangle = \sum_I \text{Pfaff}_I(z) \text{Pfaff}_{I^c}(y)$$

and in the next subsection we will see how to evaluate the right hand side in polynomial time. Then in §5 we prove that if Γ_P is planar, there is always a valid ordering.

4.3. Evaluating the complementary pairing of vectors of sub-Pfaffians. Let n be even. For an even set $I \subseteq [n]$, define $\sigma(I) = \sum_{i \in I} i$, and define $\text{sgn}(I) = (-1)^{\sigma(I) + |I|/2}$. Proofs of the following lemma can be found in [17, p. 110] and [12, p. 141].

Lemma 4.4. Let z and y be skew-symmetric $n \times n$ matrices. Then

$$\text{Pfaff}(z + y) = \sum_{\substack{I \subseteq [n], \\ |I| \equiv 0 \pmod{2}}} \text{sgn}(I) \text{Pfaff}_I(z) \text{Pfaff}_{I^c}(y)$$

To use Lemma 4.4 to compute inner products we need to adjust one of the matrices to correct the signs. For a matrix z define a matrix \tilde{z} by setting $\tilde{z}_j^i = (-1)^{i+j+1} z_j^i$. Let z be an $n \times n$ skew-symmetric matrix. Then for every even $I \subseteq [n]$,

$$\text{Pfaff}_I(\tilde{z}) = \text{sgn}(I) \text{Pfaff}_I(z).$$

For odd $|I|$, both sides are zero. For $|I| = 2p$, $p = 1, \dots, \lfloor \frac{n}{2} \rfloor$,

$$\text{Pfaff}_I(\tilde{z}) = (-1)^{i_1+i_2+1} \dots (-1)^{i_{2p-1}+i_{2p}+1} \text{Pfaff}_I(z) = \text{sgn}(I) \text{Pfaff}_I(z).$$

Thus we have the following Theorem.

Theorem 4.5. Let z, y be skew-symmetric $n \times n$ matrices. Then

$$\langle \text{sPf}(z), \text{sPf}^\vee(y) \rangle = \text{Pfaff}(\tilde{z} + y).$$

In Section 5 we show that if Γ_P is planar there is an easily computable valid edge ordering \bar{E} . Our result may be summarized as follows:

Definition 4.6. Define a problem \mathcal{P} to be *holographic-friendly* if for each instance P of \mathcal{P}

- (1) There exists a change of basis in \mathbb{C}^2 after which all the $G_i = \alpha_i \text{sPf}(x_i)$, $R_s = \beta_s \text{sPf}^\vee(y_i)$ satisfy the Grassmann-Plücker identities with complementary indexing.
- (2) There exists a valid edge order \bar{E} .

Theorem 4.7. *Let P be an instance of a holographic friendly problem \mathcal{P} . Perform the desired change of basis and write $G_i = \alpha_i \text{sPf}(x_i)$, $R_s = \beta_s \text{sPf}^\vee(y_i)$ for the generator and recognizer tensors.*

Let $\alpha = \prod_i \alpha_i$ and $\beta = \prod_s \beta_s$. Consider skew symmetric matrices x, y where x_j^i is the entry of $\alpha^{-1}G \bar{E}$ -corresponding to $I = (i, j)$ and y_j^i is the entry of $\beta^{-1}\tau(R) \bar{E}$ -corresponding to $I^c = (i, j)$. Then the number of satisfying assignments to P is given by $\alpha\beta \text{Pfaff}(\tilde{z} + y)$.

Example 4.8. Figure 2 shows an example of #Pl-Mon-3-NAE-SAT, with an edge order given by a path through the graph. The corresponding matrix, $\tilde{z} + y$ is below. In a generator order,

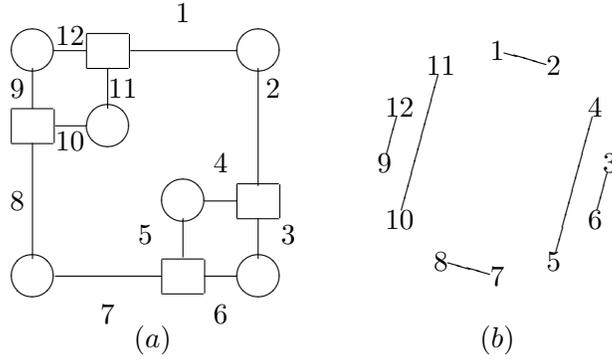


FIGURE 2. Example 4.8, and the term $S = (1, 2)(3, 6)(4, 5)(7, 8)(9, 12)(10, 11)$ in the Pfaffian (which has no crossings). Circles are variable vertices and rectangles clause vertices.

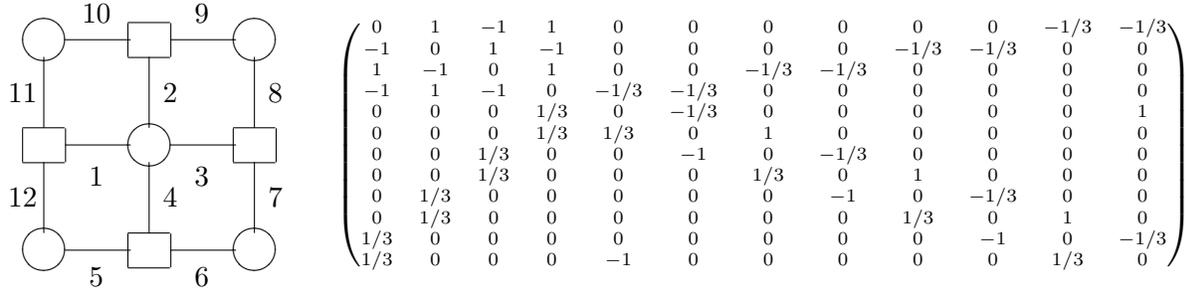
each variable corresponds to a $\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$ block. In a recognizer order, each clause corresponds to a 3×3 block with $-1/3$ above the diagonal. Sign flips $z \mapsto \tilde{z}$ occur in a checkerboard pattern with the diagonal flipped; here no flips occur. We pick up a factor of $\frac{6}{23}$ for each clause and 2 for each variable, so $\alpha = 2^6$, $\beta = (\frac{6}{23})^4$, and $\alpha\beta \text{Pfaff}(\tilde{z} + y) = 26$ satisfying assignments.

$$\tilde{z} + y = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1/3 & -1/3 \\ -1 & 0 & -1/3 & -1/3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1/3 & 0 & -1/3 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1/3 & 1/3 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & -1/3 & -1/3 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 1/3 & 0 & -1/3 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1/3 & 1/3 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & -1/3 & -1/3 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1/3 & 0 & -1/3 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1/3 & 1/3 & 0 & 1 & 0 \\ 1/3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & -1/3 \\ 1/3 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 1/3 & 0 \end{pmatrix}$$

Example 4.9. Another #Pl-Mon-3-NAE-SAT example which is not read-twice and its $\tilde{z} + y$ matrix are shown in Figure 3. The central variable has a submatrix which is again ones above the diagonal and also contributes 2 to α , so $\alpha = 2^5$, $\beta = (\frac{6}{23})^4$. Four sign changes are necessary in \tilde{z} . The result is $\alpha\beta \text{Pfaff}(\tilde{z} + y) = 14$ satisfying assignments.

5. EDGE ORDERING AND SIGN

Throughout this section we assume P is an instance of a problem \mathcal{P} where condition (1) in Definition 4.6 is satisfied, and we study when condition (2) may also be satisfied. We do not require symmetric signatures.

FIGURE 3. Another #PI-Mon-3-NAE-SAT example and its $\tilde{z} + y$ matrix.

Given an order \bar{E} we would like to know if it is valid, i.e., the signs can be made to work out right. Say \bar{E}_G, \bar{E}_R are generator and recognizer orders so that there exist skew-symmetric matrices z, y such that with respect to these orders $G = \text{sPf}(z)$, $R = \text{sPf}(y)$. Let $\pi, \tau \in \mathfrak{S}_{|E|}$ respectively be the permutations such that

$$(5) \quad \pi(\bar{E}_G) = \bar{E} \quad \text{and} \quad \tau(\bar{E}_R) = \bar{E}.$$

Then for all $J \subset [n]$, $\text{Pfaff}_{\pi(J)}(\pi(z)) = \text{sgn}(\pi|_J) \text{Pfaff}_J(z)$ and similarly for τ , so G and R are vectors of sub-Pfaffians except that possibly some signs are wrong. Valid orderings yield π, τ which preserve sub-Pfaffian signs.

We describe one type of valid ordering for planar graphs, called a *C-ordering*. For any planar bipartite graph Γ_P , a plane curve C intersecting every edge once corresponds to a non-self-intersecting Eulerian cycle in the dual of Γ_P and can be computed in $O(|E|)$ time. Fix such a C , an orientation and a starting point for C , and let \bar{E}^C be the order in which the resulting path crosses the edges of Γ_P . Define \bar{E}_G^C to be the generator order chosen so that the permutation $\pi : \bar{E}_G^C \rightarrow \bar{E}^C$ is lexicographically minimal. In particular, \bar{E}_G^C agrees with \bar{E}_G on the edges incident to any fixed generator in V . For example, the generator order on Figure 2 is 1, 2, 3, 6, 4, 5, 7, 8, 9, 12, 10, 11. Define \bar{E}_R^C similarly.

To show that \bar{E}^C is valid we will need another characterization of the sub-Pfaffians and the notion of crossing number. Let $S = \{(e_1, e'_1), \dots, (e_k, e'_k)\}$ be a partition of an ordered set I , with $|I| = 2k$, into unordered pairs. Assume, for convenience, that $e_r < e'_r$ for $1 \leq r \leq k$. Define the *crossing number* $\text{cr}(S)$ of S as

$$\text{cr}(S) = \#\{(r, s) \mid e_r < e_s < e'_r < e'_s\}.$$

Note that $\text{cr}(S)$ can be interpreted geometrically as follows. If the elements of I are arranged on a circle in order and the pairs of elements corresponding to pairs in S are joined by straight-line edges, then $\text{cr}(S)$ is the number of crossings in the resulting geometric graph (see Figure 2(b)). When the order \bar{E} on I is unclear from context we write $\text{cr}(S, \bar{E})$, instead of $\text{cr}(S)$.

For $I \subseteq E(\Gamma)$, denote by Γ_I the subgraph of Γ induced by I . Let $\mathcal{S}(\Gamma_I)$ be the set of pairings $S = \{(e_1, e'_1), \dots, (e_k, e'_k)\}$ of I such that edges in each pair share a vertex in the set V of generators. In other words, $(e_i, e'_i) \in S$ implies there exists $j \in V, s, t \in U$ such that $e_i = (j, s), e'_i = (j, t)$. In what follows we focus on generators, the corresponding statements for recognizers will be clear.

Proposition 5.1. *Let Γ be a bipartite graph and let \bar{E}_G be a generator edge order. Assume the hypotheses of Proposition 4.2 are satisfied with z the skew-symmetric $|E| \times |E|$ matrix such that $\text{sPf}(z) = G$ with the order \bar{E}_G . Let $I \subset [n] \cong E$. Then*

$$G_I = \text{Pfaff}_I(z) = \sum_{S \in \mathcal{S}(\Gamma_I)} (-1)^{\text{cr}(S)} z_S$$

where z_S is the product $\prod_{(e_i, e'_i) \in S} z_{e_i, e'_i}$.

Proof. Let $\sigma(S)$ denote the permutation

$$\sigma(S) = (e_1 \ e'_1 \ e_2 \ e'_2 \ \dots \ e_k \ e'_k).$$

By [16, p. 91] or direct verification, $\text{sgn}(\sigma(S)) = (-1)^{\text{cr}(S)}$. Therefore, for a skew-symmetric matrix z one has

$$\text{Pfaff}_I(z) = \sum_{S \in \mathcal{S}} (-1)^{\text{cr}(S)} z_S,$$

where $z_S := z_{e_1 e'_1} \dots z_{e_k e'_k}$ and the sum is taken over the set \mathcal{S} of *all* partitions of I into pairs.

We need to show that the terms $z_S, S \in \mathcal{S} \setminus \mathcal{S}(\Gamma_I)$ are zero. Note that for a nonzero term, there must be an even number of edges in the restriction to each variable. If S contains a pair with split ends $(x_i c_s, x_k c_t), i \neq k$, then $z_S = 0$. \square

The analogous statement to Proposition 5.1 holds for recognizers. We can now prove the following Lemma, which shows holographic friendly problems exist.

Lemma 5.2. *Let \mathcal{P} be a problem satisfying condition (1) of Definition 4.6. Assume furthermore that for each instance P of \mathcal{P} that Γ_P is planar. Let \bar{E}^C be a C -ordering. If π, z are defined as in (5), then $\text{sPf}(\pi(z)) = \pi(\text{sPf}(z))$.*

Proof. It suffices to show that for any $I \subseteq E(\Gamma)$ and any partition $S \in \mathcal{S}(\Gamma_I)$ of I , the signs of the term corresponding to S in $\text{Pfaff}_I(z)$ and $\text{Pfaff}_{\pi(I)}(\pi(z))$ are identical. By Proposition 5.1, this is equivalent to showing that

$$(6) \quad (-1)^{\text{cr}(S, \bar{E}^C)} = \prod_{x \in V} (-1)^{\text{cr}(S|_x, \bar{E}_G^C)},$$

where the left hand side of (6) is the sign of the term corresponding to S appearing in $\text{Pfaff}_{\pi(I)}(\pi(z))$, and the right hand side is the sign of the term corresponding to S in $\text{Pfaff}_I(z)$, as

$$\text{Pfaff}_I(z) = \prod_{x \in V} \text{Pfaff}_{I|_x}(z).$$

Here $S|_x$ and $I|_x$ denote the restriction to the edges incident to x of S and I , respectively.

A stronger equality, namely $\text{cr}(S, \bar{E}^C) = \sum_{x \in V} \text{cr}(S|_x, \bar{E}_G^C)$, holds. The curve C determining \bar{E}^C separates V from U . To exploit the geometric intuition presented above, we replace each vertex in $x \in V$ by a small circle and join the ends of edges in I on this circle by line segments corresponding to pairs in $S|_x$. The total number of crossings in the resulting graph is $\sum_{x \in V} \text{cr}(S|_x, \bar{E}_G^C) = \sum_{x \in V} \text{cr}(S|_x, \bar{E}^C)$, as \bar{E}_G^C and \bar{E}^C coincide on the set of edges incident to a fixed $x \in V$. On the other hand, a pair $\{r, s\}$ is counted in $\text{cr}(S, \bar{E}^C)$, if and only if the curves with ends on C corresponding to $e_r \cup e'_r$ and $e_s \cup e'_s$ cross.

In other words, we are considering restrictions of (the union of e_r and e'_r) and (the union e_s and e'_s) to the region of the plane bounded by C containing V . \square

It follows from Lemma 5.2 and a symmetric statement for τ that \bar{E}^C is valid.

Example 5.3. An example is given in Figure 4. There, the curves composed of edges (3 and 5), and (4 and 6) cross, and that shows that the permutation (3 5 4 6) is odd. The edges corresponding to, say, (3 and 5) and (2 and 7) don't cross, and the permutation (3 5 2 7) is even. In the example,

$$S = \{\{1, 9\}, \{2, 7\}, \{3, 5\}, \{4, 6\}, \{8, 10\}\}.$$

The term corresponding to S in $\text{Pfaff}_{\pi(I)}(\pi(x))$ is $(-1)^2 x_{1,9} x_{2,7} x_{3,5} x_{4,6} x_{8,10}$, as $\text{cr}(S, \bar{E}^C) = 2$. The term in $\text{Pfaff}(x)$ is a product of $-x_{3,5} x_{4,6}$ and $-x_{1,9} x_{2,7} x_{8,10}$, which are the terms in Pfaffians of blocks corresponding to x and x' , respectively.

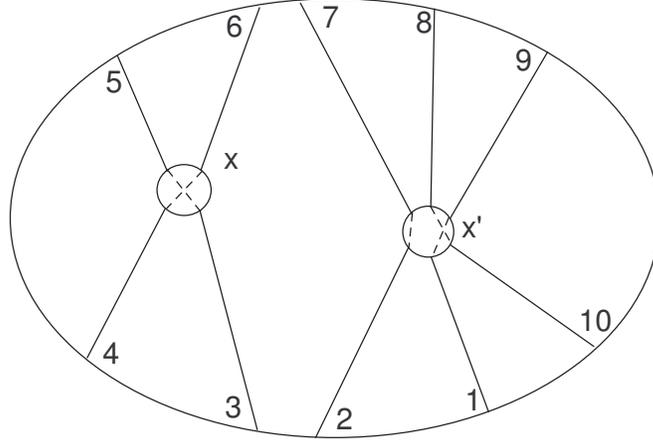


FIGURE 4. x, x' are two generators, the oval is C and the numbers indicate the ordering of the edges determined by C

6. BEYOND PFAFFIANS

As mentioned in the introduction, the key to our approach is that the pairing of a vector in a vector space of dimension 2^n with a vector in its dual space can be accomplished by evaluating a Pfaffian if both vectors are vectors of Pfaffians of some skew-symmetric matrix. This type of simplification occurs in other situations as explained in Appendix §7 below. One simple case is that if the vector space is of dimension $\binom{n}{k}$ and the vectors that are to be paired are vectors of minors of some $k \times (n - k)$ matrix. Then the pairing can be done by computing the determinant of an easily constructed auxiliary $(n - k \times n - k)$ or $(k \times k)$ -matrix, so if k is on the order of $\lfloor \frac{n}{2} \rfloor$ there is a spectacular savings. Explicitly, for $k \times \ell$ matrices z and y , with $G = \text{sDet}(z)$ and $R = \text{sDet}(y)$,

$$\langle G, R \rangle = \det(\text{Id} + z^\top y).$$

Here is an example that exploits this situation.

Example 6.1. The following result of [10] can be interpreted as a holographic algorithm. Given a graph G and an arbitrary orientation of $E(G)$, the *incidence matrix* $B = (b_v^e)_{v \in V(G), e \in E(G)}$ is a $|V(G)| \times |E(G)|$ matrix defined by

$$b_v^e = \begin{cases} 1 & \text{if } v \text{ is the initial vertex of } e, \\ -1 & \text{if } v \text{ is the terminal vertex of } e, \\ 0 & \text{otherwise.} \end{cases}$$

For $W \subseteq V(G)$ and $F \subseteq E(G)$, with $|W| = |F|$, let $\Delta_{W,F}(B)$ denote the corresponding minor of B . Let $\text{sDet}(B) = (1, \Delta_{v,e}B, \dots, \Delta_{W,F}(B), \dots)$ denote the vector of minors of B .

A *rooted spanning forest* of G is a pair (H, W) , where $W \subseteq V(G)$, H is a spanning acyclic subgraph of G , and every component of H contains exactly one vertex of W . The minor $\Delta_{W,F}(B)$ equals to ± 1 if $(G|_F, V(G) - W)$ is a rooted spanning forest, and $\Delta_{W,F}(B) = 0$, otherwise. (See [10] for a proof of a generalization of this statement to weighted graphs.) Therefore, the

value of the pairing

$$\langle \text{sDet}(B^t), \text{sDet}(B) \rangle = \sum_{W \subseteq V(G)} \sum_{F \subseteq V(G)} (\Delta_{W,F}(B))^2$$

is equal to the number of rooted spanning forests of G . It is shown [10] that this value can be computed efficiently by the Cauchy-Binet formula:

$$\sum_{W \subseteq V(G)} \sum_{F \subseteq V(G)} (\Delta_{W,F}(B))^2 = \det(\text{Id} + B^t B),$$

where Id is a $|E(G)| \times |E(G)|$ identity matrix.

From our point of view the result outlined in this example is an instance of the above fact that the pairing of vectors in the Grassmannian and its dual can be computed efficiently. (The Grassmannian can be locally parametrized by vectors of minors of matrices.) The above efficient algorithm for counting rooted spanning forests is surprising in the same sense as many holographic algorithms are: A closely related problem of enumerating spanning forests of a graph is $\#P$ -hard [13].

ACKNOWLEDGMENTS

This paper is an outgrowth of the AIM workshop *Geometry and representation theory of tensors for computer science, statistics and other areas* July 21-25, 2008, and authors gratefully thank AIM and the other participants of the workshop. We especially thank J. Cai, P. Lu, and L. Valiant for their significant efforts to explain their theory to us during the workshop. J. Cai is also to be thanked for continuing to answer our questions with extraordinary patience for months afterwards. We thank R. Thomas for his help with the graph-theoretical part of the argument.

REFERENCES

1. Jin-Yi Cai and Vinay Choudhary, *Some results on matchgates and holographic algorithms*, Automata, languages and programming. Part I, Lecture Notes in Comput. Sci., vol. 4051, Springer, Berlin, 2006, pp. 703–714.
2. ———, *Valiant’s holant theorem and matchgate tensors*, Theory and applications of models of computation, Lecture Notes in Comput. Sci., vol. 3959, Springer, Berlin, 2006, pp. 248–261.
3. ———, *Valiant’s holant theorem and matchgate tensors*, Theoret. Comput. Sci. **384** (2007), no. 1, 22–32.
4. Jin-Yi Cai and Pinyan Lu, *Holographic algorithms: from art to science*, STOC’07—Proceedings of the 39th Annual ACM Symposium on Theory of Computing, ACM, New York, 2007, pp. 401–410.
5. ———, *Holographic algorithms: the power of dimensionality resolved*, Automata, languages and programming, Lecture Notes in Comput. Sci., vol. 4596, Springer, Berlin, 2007, pp. 631–642.
6. ———, *On symmetric signatures in holographic algorithms*, STACS 2007, Lecture Notes in Comput. Sci., vol. 4393, Springer, Berlin, 2007, pp. 429–440.
7. ———, *Basis collapse in holographic algorithms*, Comput. Complexity **17** (2008), no. 2, 254–281.
8. Jin-Yi Cai, Pinyan Lu, and Mingji Xia, *Holographic algorithms by fibonacci gates and holographic reductions for hardness*, Proceedings of the 49th annual Symposium on Foundations of Computer Science (2008), 644–653.
9. Claude Chevalley, *The algebraic theory of spinors and Clifford algebras*, Springer-Verlag, Berlin, 1997, Collected works. Vol. 2, Edited and with a foreword by Pierre Cartier and Catherine Chevalley, With a postface by J.-P. Bourguignon.
10. F. R. K. Chung and Robert P. Langlands, *A combinatorial Laplacian with vertex weights*, J. Combin. Theory Ser. A **75** (1996), no. 2, 316–327.
11. B. Hayes, *Accidental algorithms*, American Scientist **96** (2008), no. 1, 9–13.
12. Masao Ishikawa and Masato Wakayama, *Applications of minor-summation formula. II. Pfaffians and Schur polynomials*, J. Combin. Theory Ser. A **88** (1999), no. 1, 136–157.

13. F. Jaeger, D. L. Vertigan, and D. J. A. Welsh, *On the computational complexity of the Jones and Tutte polynomials*, Math. Proc. Cambridge Philos. Soc. **108** (1990), no. 1, 35–53.
14. P. W. Kasteleyn, *Graph theory and crystal physics*, Graph Theory and Theoretical Physics, Academic Press, London, 1967, pp. 43–110.
15. Jason Morton, *Pfaffian circuits*, preprint arXiv:1101.0129.
16. Serguei Norine, *Pfaffian graphs, T-joins and crossing numbers*, Combinatorica **28** (2008), no. 1, 89–98.
17. John R. Stembridge, *Nonintersecting paths, Pfaffians, and plane partitions*, Adv. Math. **83** (1990), no. 1, 96–131.
18. H. N. V. Temperley and Michael E. Fisher, *Dimer problem in statistical mechanics—an exact result*, Philos. Mag. (8) **6** (1961), 1061–1063.
19. Leslie G. Valiant, *Quantum computers that can be simulated classically in polynomial time*, Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing (New York), ACM, 2001, pp. 114–123 (electronic).
20. ———, *Expressiveness of matchgates*, Theoret. Comput. Sci. **289** (2002), no. 1, 457–471.
21. ———, *Quantum circuits that can be simulated classically in polynomial time*, SIAM J. Comput. **31** (2002), no. 4, 1229–1254.
22. ———, *Accidental algorithms*, In Proc. 47th Annual IEEE Symposium on Foundations of Computer Science 2006, 2004, pp. 509–517.
23. ———, *Holographic algorithms (extended abstract)*, Proceedings of the 45th annual Symposium on Foundations of Computer Science (2004), 306–315.
24. ———, *Holographic circuits*, Automata, languages and programming, Lecture Notes in Comput. Sci., vol. 3580, Springer, Berlin, 2005, pp. 1–15.
25. ———, *Holographic algorithms*, SIAM J. Comput. **37** (2008), no. 5, 1565–1594.
26. Leslie G. Valiant and V.V. Vazirani, *NP is as easy as detecting unique solutions*, Theoretical Computer Science **47** (1986), 85–93.

7. APPENDIX: SPINORS AND HOLOGRAPHIC ALGORITHMS

The Grassmann-Plücker identities are the defining equations for the *spinor varieties* (set of pure spinors). These equations date back to Chevalley in the 1950’s [9]. The spinor varieties, of which there are two (isomorphic to each other) for each n , $\hat{\mathcal{S}}_+$, $\hat{\mathcal{S}}_-$, respectively live in $\Lambda^{even}\mathbb{C}^n =: \mathcal{S}_+$, and $\Lambda^{odd}\mathbb{C}^n =: \mathcal{S}_-$. The parity condition corresponds to requiring that G, R both be either in \mathcal{S}_+ or \mathcal{S}_- . If n is odd then $\mathcal{S}_+, \mathcal{S}_-$ are dual vector spaces to one another, and if n is even, each is self-dual. It is this self-duality that leads to the simplification of the exposition with n is even - the discussion for n odd is given below.

The spinor varieties admit a cover by Zariski open subsets where each subset in e.g. $\hat{\mathcal{S}}_+$ is covered by a map of the form

$$(7) \quad \phi : \Lambda^2\mathbb{C}^n \rightarrow \bigoplus_j \Lambda^{2j}\mathbb{C}^n = \Lambda^{even}\mathbb{C}^n$$

$$(8) \quad x \mapsto \text{sPf}(x)$$

where for each $I \subset \{1, \dots, n\}$ of even cardinality $\text{Pfaff}_I(c) \in \Lambda^{|I|}\mathbb{C}^n$.

The identification $\mathcal{S}_+ \simeq \Lambda^{even}\mathbb{C}^n$ is not canonical. We obtain different identifications by composing ϕ with the action of the Weyl group $\mathfrak{S}_n \times \mathbb{Z}_2$. The Weyl group action assures that some “less convenient” map will have first entry nonzero for G, R as mentioned in §4.2.

The map (7) is a special case of a natural map to the “big cell” in a compact Hermitian symmetric space and the potential generalizations to holographic algorithms mentioned in the introduction would correspond to replacing $\hat{\mathcal{S}}_+$ by a Lagrangian Grassmannian or an ordinary Grassmannian of k -planes in a n -dimensional space. More generally, if V is a generalized $G(n)$ -cominuscule module, where n denotes the rank of the semi-simple group G , then the pairing $V \times V^* \rightarrow \mathbb{C}$, when restricted to the cone over the closed orbits in V, V^* can be computed with $O(n^4)$ arithmetic operations, even though the dimension of V is generally exponential in n .

Much of the exposition could be rephrased more concisely using the language of representation theory. For example, the fact that if each G_i lies in a small spinor variety then $G = \otimes G_i$ lies

in a spinor variety as well, is a consequence that the tensor product of highest weight vectors of subgroups with compatible Weyl chambers will be a highest weight vector for the larger group. Similarly the map $z \mapsto \tilde{z}$ has a natural interpretation in terms of an involution on the Clifford module structure that \mathcal{S}_+ comes equipped with.

On the other hand \mathbb{C}^{2^n} may be viewed as $(\mathbb{C}^2)^{\otimes n}$ and as such, inherits an $SL_2\mathbb{C}$ -action. The $SL_2(\mathbb{C})$ action corresponds to the changes of basis, and what we are trying to do is determine which pairs of points can by simultaneously be moved into the spinor varieties in $(\mathbb{C}^2)^{\otimes n}$ and the dual space $(\mathbb{C}^{2*})^{\otimes n}$. The convenient basis referred to in the text corresponds to an identification that embeds the torus of SL_2 diagonally into the torus of $Spin_{2n}$ so weight vectors map to weight vectors.

To continue the group perspective in complexity theory more generally, one can also view the ability to compute the determinant quickly via Gaussian elimination as the consequence of the robustness of the action of the group preserving the determinant: whereas above there is a subvariety of a huge space (the spinor variety) on which the pairing can be computed quickly, and a group SL_2 that preserves the pairing - a holographic algorithm can be exploited if the pair (G, R) can be moved into the subvariety $\hat{\mathbb{S}}_+ \times \hat{\mathbb{S}}_+$ under the action of SL_2 . In Gaussian elimination, for the corresponding subvariety one takes, e.g., the set of upper-triangular matrices, and the group preserving the determinant acts on the space of matrices sufficiently robustly that any matrix can be moved into this subvariety (and in polynomial time). Contrast this with the permanent which is also easy to evaluate on upper-triangular matrices, but the group preserving the permanent is not sufficiently robust to send an arbitrary matrix to an upper-triangular one. This difference in robustness of group actions might explain the difference between the determinant and permanent, as well as why only solutions to certain *SAT* problems can (so far) be counted quickly.

8. APPENDIX: NON-SYMMETRIC SIGNATURES

Most of the natural examples of holographic algorithms, and, in particular, the examples given in this paper, correspond to generator and recognizer signatures G_i and R_s which are *symmetric*, that is invariant under permutations of edges incident to the corresponding vertex. The assumption that the signatures are symmetric is also convenient for our arguments. If the signatures are symmetric, then the generator tensor G can be represented as a vector of sub-Pfaffians in some generator order if and only if it can be represented as such a vector in every generator order, and the same holds for recognizer orders. This does not hold for general, non-symmetric signatures. We now explain how to deal with non-symmetric signatures.

It is shown in Section 5 that given a planar curve C , an edge order \bar{E}^C and a generator order \bar{E}_G^C , the tensor G can be represented as a vector of sub-Pfaffians in \bar{E}^C if and only if it can be represented as one in \bar{E}_G^C . A similar statement holds for \bar{E}^C and the recognizer order \bar{E}_R^C . The edges incident to a given generator are ordered in a clockwise cyclic order in \bar{E}_G^C . It is easy to verify that only a cyclic, not linear, ordering enters Grassmann-Plücker identities. Thus for non-symmetric signatures the following statement holds.

Theorem 8.1. *Let \mathcal{P} be a problem such that each instance P of \mathcal{P} admits a matchgate formulation $\Gamma_P = (V, U, E)$ with Γ_P planar. Let the edges incident to every vertex of Γ_P be ordered in a clockwise order. Assume condition (1) of Definition 4.6 holds. Then there exists a valid order, i.e. condition (2) is also satisfied and thus the number of satisfying assignments of P can be found in polynomial time.*

The assumption that the edges are ordered in a way that agrees with a planar embedding is also used in the matchgate formulation, as the matchgates must be inserted in such a way that the resulting graph remains planar.