Electronic Colloquium on Computational Complexity, Report No. 337 (2026)

Pfaffian Recognition is in coRP

Priyanshu Pant
Indian Institute of Technology Indore
priyanshupant03@gmail.com

Ranveer Singh
Indian Institute of Technology Indore
ranveer@iiti.ac.in

Abstract

A graph is called Pfaffian if it admits a Pfaffian orientation. The Pfaffian recognition problem asks whether a given graph is Pfaffian. In 1989, Vazirani and Yannakakis showed that this problem is polynomial time equivalent to testing whether a prescribed orientation is Pfaffian, and that the problem belongs to coNP. The complexity of the problem for general graphs has remained open since then. For bipartite graphs, polynomial time recognition algorithms were obtained by Robertson, Seymour, and Thomas in 1999, and independently by McCuaig in 2004. We show that testing whether a given orientation is Pfaffian can be reduced to polynomial identity testing. This gives a randomized polynomial time algorithm for Pfaffian recognition on general graphs, placing the problem in coRP.

1 Introduction

Counting perfect matchings is a central problem in combinatorics and the theory of counting problems [6, 8]. A *perfect matching* of a graph is a set of pairwise disjoint edges that covers every vertex exactly once. Although the existence of a perfect matching can be decided in polynomial time [3], counting the number of perfect matchings is much harder. Valiant proved that the problem is #P-complete, even for bipartite graphs [17].

All graphs considered in this paper are finite and simple. For a graph G, we write $V(G)$ and $E(G)$ for its vertex and edge sets, and $\mathcal{M}(G)$ for its set of perfect matchings. Let G have $2n$ vertices, and assign a variable x_e to each edge $e \in E(G)$. The *perfect matching polynomial* of G is

$$\mathrm{PerfMatch}(G) = \sum_{M \in \mathcal{M}(G)} \prod_{e \in M} x_e.$$

Setting all variables $x_e = 1$ gives the number of perfect matchings of G.

A classical algebraic approach to counting perfect matchings uses Pfaffian orientations. An *orientation* D of G is a choice of direction for each edge. Fix an ordering $1, 2, \ldots, 2n$ of the vertices of G. If $e = ij \in E(G)$, we write x_e and x_{ij} interchangeably for the corresponding edge variable. The *weighted skew adjacency matrix* of G with respect to D, denoted by A_D, is the $2n \times 2n$ matrix whose (i, j)-entry is

$$(A_D)_{ij} = \begin{cases} x_{ij}, & \text{if } ij \in E(G) \text{ is oriented from } i \text{ to } j, \\ -x_{ij}, & \text{if } ij \in E(G) \text{ is oriented from } j \text{ to } i, \\ 0, & \text{if } ij \notin E(G). \end{cases}$$

Let $M = \{e_1, e_2, \ldots, e_n\}$ be a perfect matching of G, and suppose that in the orientation D, the tail and the head of e_i are u_i and v_i, respectively. The *permutation of M in D*, denoted by $\pi_D(M)$, is

$$\pi_D(M) = \begin{pmatrix} 1 & 2 & 3 & 4 & \cdots & 2n-1 & 2n \\ u_1 & v_1 & u_2 & v_2 & \cdots & u_n & v_n \end{pmatrix}.$$

ISSN 1433-8092

The *sign of M with respect to D*, denoted by $\text{sign}_D(M)$, is the sign of $\pi_D(M)$. This value is independent of the order in which the edges of M are enumerated. The Pfaffian of A_D is

$$\text{Pf}(A_D) = \sum_{M \in \mathcal{M}(G)} \text{sign}_D(M) \prod_{e \in M} x_e.$$

It contains exactly the same monomials as $\text{PerfMatch}(G)$, with signs determined by the orientation. If all perfect matchings have the same sign with respect to D, then

$$\text{Pf}(A_D) = \pm \text{PerfMatch}(G).$$

Such an orientation is called a *Pfaffian orientation*, and a graph admitting one is called *Pfaffian*. Once a Pfaffian orientation is known, the number of perfect matchings can be computed in polynomial time using Cayley's identity [1],

$$\det(A) = \text{Pf}(A)^2,$$

which holds for every skew symmetric matrix A of even order.

Pfaffian orientations originate in Kasteleyn's work on the dimer problem. In 1967, Kasteleyn showed that every planar graph admits a Pfaffian orientation and gave a polynomial time algorithm for finding one, yielding a polynomial time algorithm for counting perfect matchings in planar graphs [6]. In 1975, Little [7] proved that a bipartite graph with a perfect matching is Pfaffian if and only if it contains no even subdivision of $K_{3,3}$ whose deletion leaves a graph with a perfect matching. In 1999, Robertson, Seymour, and Thomas [14], and independently McCuaig in 2004 [10], obtained structural characterizations which lead to polynomial time algorithms for deciding whether a bipartite graph is Pfaffian.

Considerable progress has also been made for some nonbipartite graph classes. Fischer and Little [4] characterized Pfaffian near-bipartite graphs, which are nonbipartite matching-covered graphs that become bipartite and remain matching covered after deleting two specific edges. Miranda and Lucchesi [11] later gave a polynomial time Pfaffian recognition algorithm for weakly near-bipartite graphs, a larger class that includes near-bipartite graphs. Carvalho, Lucchesi, and Murty [2] extended Little's characterization to solid matching-covered graphs, a class that includes bipartite matching-covered graphs.

For general graphs, however, the problem has proved much more difficult. Norine and Thomas [13] constructed an infinite family of minimal non-Pfaffian graphs with respect to matching minors, showing that the forbidden obstruction characterization for bipartite graphs does not extend directly to general graphs. Despite the progress on special graph classes, the existence of a polynomial time algorithm for recognizing general Pfaffian graphs remained open; see [16, 9].

A major insight into the Pfaffian recognition problem came from Vazirani and Yannakakis in 1989 [18]. Given a graph G, they construct in polynomial time an orientation $D^\star$ such that G admits a Pfaffian orientation if and only if $D^\star$ is Pfaffian. Thus, the recognition problem reduces to testing whether this one prescribed orientation $D^\star$ is Pfaffian. Using this reduction, they also showed that Pfaffian graph recognition belongs to coNP. Recall that a decision problem is in coNP if every no instance has a polynomial-size certificate that can be verified in polynomial time. If two perfect matchings have different signs with respect to a given orientation, they certify that the orientation is not Pfaffian. Such a pair can be checked in polynomial time. Hence, when G is not Pfaffian, two perfect matchings of different signs in $D^\star$ give such a certificate.

However, testing whether a given orientation is Pfaffian is not straightforward, since a graph may have exponentially many perfect matchings. In this paper, we give a randomized polynomial time algorithm for this problem using an algebraic approach, establishing the following result.

Theorem 1.1. *Pfaffian graph recognition is in* coRP.

Recall that a decision problem belongs to coRP if there is a randomized polynomial time algorithm that always accepts yes instances and rejects no instances with probability at least 1/2. For Pfaffian recognition, given a graph G, we first use the reduction of Vazirani and Yannakakis to construct the orientation $D^\star$. We then run our randomized test on $D^\star$. If G is Pfaffian, then $D^\star$ is Pfaffian and our test always accepts. If G is not Pfaffian, then $D^\star$ is non-Pfaffian and our test rejects with probability at least 1/2. Repeating the test independently reduces the probability of error exponentially.

The proof reduces testing a prescribed orientation to polynomial identity testing. Given G and an orientation D, we define a polynomial $P_D(\mathbf{x})$ such that

$$D \text{ is Pfaffian} \quad \Longleftrightarrow \quad P_D(\mathbf{x}) \equiv 0.$$

The polynomial P_D records the perfect matchings whose sign is opposite to that of a fixed perfect matching. We show how to evaluate P_D in polynomial time using determinants of principal submatrices of the Tutte matrix. We further show that one randomized trial for testing a given orientation can be implemented using $O(|V(G)|^3)$ field operations.

The rest of the paper is organized as follows. In Section 2 we study the signs of perfect matchings and define the polynomial used in the algorithm. In Section 3 we use this polynomial to test whether an orientation is Pfaffian and thus obtain the recognition algorithm. In Section 4 we give a faster way to evaluate the polynomial and obtain a randomized test for a given orientation using $O(|V(G)|^3)$ field operations.

2 A polynomial characterization of Pfaffianness

In this section, we construct a polynomial that is identically zero exactly when the prescribed orientation is Pfaffian.

2.1 Signs of perfect matchings

We use $\mathbb{F}_2 = \{0, 1\}$ to record parity, with 0 representing even parity and 1 representing odd parity. For an edge $e = ij$, with $i < j$, say that e is *backward* in D if it is oriented from j to i. For a perfect matching M, let $\mathrm{rev}_D(M) \in \mathbb{F}_2$ be the parity of the number of backward edges in M.

Two vertex disjoint edges $\{i, j\}$ and $\{k, \ell\}$, written with $i < j$ and $k < \ell$, *cross* if their endpoints alternate, that is, if $i < k < j < \ell$ or $k < i < \ell < j$. For a perfect matching M, let $\mathrm{cr}(M) \in \mathbb{F}_2$ be the parity of the number of crossing pairs in M. The following sign formula also appears in Norine [12].

Lemma 2.1. *For every perfect matching M,*

$$\mathrm{sign}_D(M) = (-1)^{\mathrm{cr}(M)+\mathrm{rev}_D(M)}.$$

Proof. For each edge of M, write its endpoints as (a_i, b_i) with $a_i < b_i$, and order the edges so that $a_1 < a_2 < \cdots < a_n$. Consider the permutation

$$\pi = \begin{pmatrix} 1 & 2 & \cdots & 2n-1 & 2n \\ a_1 & b_1 & \cdots & a_n & b_n \end{pmatrix}.$$

Fix $i < j$. Since $a_i < a_j$, the only possible inversions involving the endpoints of the two edges $\{a_i, b_i\}$ and $\{a_j, b_j\}$ come from b_i. If $b_i < a_j$, there are no such inversions. If $a_j < b_i < b_j$, there is

exactly one inversion, while if $b_j < b_i$, there are two. Thus the number of inversions contributed by this pair of edges is odd exactly when

$$a_i < a_j < b_i < b_j,$$

which is exactly when the two edges cross. Therefore the parity of the total number of inversions of π is $\mathrm{cr}(M)$, and hence

$$\mathrm{sign}(\pi) = (-1)^{\mathrm{cr}(M)}.$$

To obtain the oriented order defining $\mathrm{sign}_D(M)$, we swap the two endpoints of each backward edge. Each such swap changes the sign by a factor of -1. Therefore

$$\mathrm{sign}_D(M) = (-1)^{\mathrm{cr}(M)+\mathrm{rev}_D(M)}. \qquad \square$$

2.2 The wrong sign polynomial

Assume that G has a perfect matching, and fix one such matching F. Put $s = \mathrm{cr}(F) + \mathrm{rev}_D(F) \in \mathbb{F}_2$. By Lemma 2.1, a perfect matching M has sign opposite to F exactly when

$$s + \mathrm{cr}(M) + \mathrm{rev}_D(M) = 1 \qquad \text{in } \mathbb{F}_2.$$

Let $\mathbf{x} = (x_e)_{e\in E(G)}$, where the x_e are variables over a finite field of characteristic two. Define

$$P_D(\mathbf{x}) = \sum_{\substack{M\in\mathcal{M}(G)\\ \mathrm{sign}_D(M)\neq\mathrm{sign}_D(F)}} \prod_{e\in M} x_e^2.$$

The polynomial depends on the choice of F. Since we fix F throughout, we write P_D without mentioning F explicitly.

Proposition 2.2. *The orientation D is Pfaffian if and only if $P_D(\mathbf{x}) \equiv 0$.*

Proof. The monomials of P_D are in one to one correspondence with the perfect matchings whose sign differs from that of F. Since F is itself a perfect matching, P_D is identically zero exactly when every perfect matching has the same sign as F, which is exactly when D is Pfaffian. $\square$

To obtain an efficient expression for P_D, we use the Tutte matrix. The *Tutte matrix* of G, denoted by $T_G(\mathbf{x})$, is the $2n \times 2n$ skew symmetric matrix defined by

$$(T_G(\mathbf{x}))_{ij} = \begin{cases} x_{ij}, & \text{if } i < j \text{ and } ij \in E(G),\\ -x_{ji}, & \text{if } i > j \text{ and } ij \in E(G),\\ 0, & \text{otherwise.}\end{cases}$$

Write $T = T_G(\mathbf{x})$. For an edge $e = ij$, let $T - e$ denote the principal submatrix obtained by deleting rows and columns i, j. For two vertex disjoint edges $e = ij$ and $f = k\ell$, let $T - e - f$ denote the principal submatrix obtained by deleting rows and columns i, j, k, ℓ.

Let $R_D \subseteq E(G)$ be the set of backward edges of D, and let $\mathcal{C}(G)$ be the set of unordered pairs $\{e, f\}$ of vertex disjoint edges which cross in the fixed vertex ordering. The next identity expresses P_D in terms of principal minors of the Tutte matrix.

Lemma 2.3. *Over every field of characteristic two,*

$$P_D(\mathbf{x}) = s\det(T) + \sum_{e\in R_D} x_e^2 \det(T - e) + \sum_{\{e,f\}\in\mathcal{C}(G)} x_e^2 x_f^2 \det(T - e - f). \tag{1}$$

Proof. In characteristic two, $-1 = 1$, so all terms in the Pfaffian expansion have positive signs. Thus

$$\mathrm{Pf}(T) = \sum_{M \in \mathcal{M}(G)} \prod_{e \in M} x_e.$$

In characteristic two, squaring both sides gives

$$(\mathrm{Pf}(T))^2 = \sum_{M \in \mathcal{M}(G)} \prod_{e \in M} x_e^2.$$

By Cayley's identity, $(\mathrm{Pf}\, T)^2 = \det T$, and hence

$$\det(T) = \sum_{M \in \mathcal{M}(G)} \prod_{e \in M} x_e^2.$$

The same argument applies to principal submatrices. Let $e = ij$, and let $G - \{i, j\}$ denote the graph obtained by deleting the vertices i and j. Since $T - e$ is the Tutte matrix of $G - \{i, j\}$,

$$\det(T - e) = \sum_{M \in \mathcal{M}(G - \{i,j\})} \prod_{a \in M} x_a^2.$$

Every perfect matching of $G - \{i, j\}$ extends uniquely to a perfect matching of G containing e, so

$$x_e^2 \det(T - e) = \sum_{\substack{M \in \mathcal{M}(G) \\ e \in M}} \prod_{a \in M} x_a^2.$$

Likewise, for two vertex disjoint edges $e = ij$ and $f = k\ell$,

$$x_e^2 x_f^2 \det(T - e - f) = \sum_{\substack{M \in \mathcal{M}(G) \\ e, f \in M}} \prod_{a \in M} x_a^2.$$

Fix a perfect matching M and consider the coefficient of $\prod_{e \in M} x_e^2$ on the right hand side of (1). The term $s \det T$ contributes s. The second sum contributes once for each backward edge in M, and the third once for each crossing pair in M. Its coefficient is therefore

$$s + \mathrm{rev}_D(M) + \mathrm{cr}(M).$$

By Lemma 2.1, this equals 1 exactly when $\mathrm{sign}_D(M) \neq \mathrm{sign}_D(F)$, and equals 0 otherwise. This is precisely the coefficient of the same monomial in $P_D(\mathbf{x})$. □

3 Randomized Pfaffian recognition

In this section, we use the determinant identity to evaluate P_D in polynomial time and, together with polynomial identity testing, obtain a randomized test for a prescribed orientation.

3.1 Testing Pfaffianness of an orientation

We use the following standard form of the Schwartz-Zippel lemma [15, 19].

Lemma 3.1 (Schwartz-Zippel)**.** *Let $P \in \mathbb{F}[x_1, \ldots, x_m]$ be a nonzero polynomial of total degree at most d, and let $S \subseteq \mathbb{F}$ be finite. If $\alpha_1, \ldots, \alpha_m$ are chosen independently and uniformly from S, then*

$$\Pr[P(\alpha_1, \ldots, \alpha_m) = 0] \leq \frac{d}{|S|}.$$

Theorem 3.2. *There is a randomized polynomial time algorithm which, given a graph G and an orientation D, always accepts if D is Pfaffian and rejects with probability at least $1/2$ if D is non-Pfaffian.*

Proof. First check whether G has a perfect matching. If it has none, then D is trivially Pfaffian, so we accept. Otherwise, find a perfect matching F and consider the polynomial $P_D(\mathbf{x})$ defined above.

Choose a field $\mathbb{F}_q$ of characteristic two, where $q = 2^t$ is the smallest power of 2 satisfying $q \geq 4n$. Then $4n \leq q < 8n$, and since $t = O(\log n)$, the basic field operations in $\mathbb{F}_q$ can be carried out in polynomial time. The randomized algorithm proceeds as follows.

1. For each edge $e \in E(G)$, choose $\alpha_e \in \mathbb{F}_q$ independently and uniformly at random.
2. Evaluate $P_D(\alpha)$ using Lemma 2.3.
3. Reject if $P_D(\alpha) \neq 0$, and accept otherwise.

If D is Pfaffian, then $P_D(\mathbf{x}) \equiv 0$ by Proposition 2.2, so the algorithm always accepts. If D is non-Pfaffian, then $P_D(\mathbf{x})$ is a nonzero polynomial of degree $2n$. By Lemma 3.1,

$$\Pr[P_D(\alpha) = 0] \leq \frac{2n}{q} \leq \frac{1}{2}.$$

A non-Pfaffian orientation is therefore rejected with probability at least $1/2$. By Lemma 2.3, evaluating $P_D(\alpha)$ requires only polynomially many determinant computations on matrices of order at most $2n$, so the algorithm runs in polynomial time. □

Repeating the test independently r times reduces the error probability exponentially. If D is Pfaffian, every trial accepts, so the repeated test always accepts. If D is non-Pfaffian, each trial accepts with probability at most $1/2$. Therefore the probability that all r independent trials accept is at most $\left(\frac{1}{2}\right)^r$. Thus, by repeating the test, the probability of incorrectly accepting a non-Pfaffian orientation can be made arbitrarily small.

3.2 Pfaffian graph recognition

We use the reduction of Vazirani and Yannakakis [18, Theorem 3.1]. Given a graph G with a perfect matching, their reduction constructs in polynomial time an orientation $D^\star$ such that G admits a Pfaffian orientation if and only if $D^\star$ is Pfaffian.

Proof of Theorem 1.1. If G has no perfect matching, then G is trivially Pfaffian. Otherwise, construct $D^\star$ using the reduction of Vazirani and Yannakakis and apply Theorem 3.2. If G admits a Pfaffian orientation, then $D^\star$ is Pfaffian, so the algorithm always accepts. If G does not admit a Pfaffian orientation, then $D^\star$ is non-Pfaffian and is rejected with probability at least $1/2$. □

4 Efficient evaluation of the wrong sign polynomial

We now show that whenever T_α is invertible, $P_D(\alpha)$ can be evaluated using $O(|V(G)|^3)$ field operations. The main idea is to compute the required principal minors from a single matrix inverse.

We use Jacobi's complementary minor identity [5]. For an invertible matrix A indexed by V and a subset $X \subseteq V$,

$$\det(A[V \setminus X]) = \det(A)\det(A^{-1}[X]),$$

where $A[X]$ denotes the principal submatrix of A with rows and columns indexed by X.

Let $\alpha = (\alpha_e)_{e\in E(G)}$ be a point over a field of characteristic two, and write $T_\alpha = T_G(\alpha)$. Set $\alpha_{ij} = 0$ whenever $ij \notin E(G)$. Suppose that T_α is invertible, and write $T_\alpha^{-1} = (t_{ij})_{i,j\in V(G)}$.

We first note that both T_α and its inverse T_α^{-1} are alternating. Over a field of characteristic two, an *alternating matrix* is a symmetric matrix with zero diagonal. For any alternating matrix $A = (a_{ij})_{n\times n}$ and any vector $v \in \mathbb{F}^n$, we have

$$v^T A v = \sum_{i,j=1}^{n} a_{ij}v_i v_j = \sum_{i=1}^{n} a_{ii}v_i^2 + \sum_{1\le i<j\le n} (a_{ij}+a_{ji})v_i v_j = 0,$$

since $a_{ii} = 0$ and $a_{ij} = -a_{ji}$. By definition, T_α is alternating. In characteristic two, symmetric and skew-symmetric matrices are the same, so $T_\alpha^T = T_\alpha$. Hence $(T_\alpha^{-1})^T = (T_\alpha^T)^{-1} = T_\alpha^{-1}$. Thus T_α^{-1} is also symmetric. To see that T_α^{-1} has zero diagonal, let e_i be the vector of length $|V(G)|$ with 1 in position i and 0 elsewhere. Then

$$t_{ii} = e_i^T T_\alpha^{-1} e_i = (T_\alpha^{-1}e_i)^T T_\alpha (T_\alpha^{-1}e_i) = 0,$$

where the last equality follows since T_α is alternating. Hence T_α^{-1} has zero diagonal and is therefore alternating. For an edge $e = ij$, Jacobi's identity gives

$$\det(T_\alpha - e) = \det(T_\alpha)\det\big(T_\alpha^{-1}[\{i,j\}]\big) = \det(T_\alpha)\det\begin{pmatrix} 0 & t_{ij} \\ t_{ij} & 0 \end{pmatrix} = \det(T_\alpha)t_{ij}^2. \tag{2}$$

Similarly, for two vertex disjoint edges $e = ij$ and $f = k\ell$, Jacobi's identity gives

$$\det(T_\alpha - e - f) = \det(T_\alpha)\det\begin{pmatrix} 0 & t_{ij} & t_{ik} & t_{i\ell} \\ t_{ij} & 0 & t_{jk} & t_{j\ell} \\ t_{ik} & t_{jk} & 0 & t_{k\ell} \\ t_{i\ell} & t_{j\ell} & t_{k\ell} & 0 \end{pmatrix} = \det(T_\alpha)\big(t_{ij}t_{k\ell} + t_{ik}t_{j\ell} + t_{i\ell}t_{jk}\big)^2. \tag{3}$$

The last equality follows since the determinant of an alternating matrix of even order is the square of its Pfaffian. Define

$$Q_D(\alpha) = s + \sum_{ij\in R_D} \alpha_{ij}t_{ij} + \sum_{\{ij,k\ell\}\in\mathcal{C}(G)} \alpha_{ij}\alpha_{k\ell}\big(t_{ij}t_{k\ell} + t_{ik}t_{j\ell} + t_{i\ell}t_{jk}\big). \tag{4}$$

Lemma 4.1. *If T_α is invertible, then $P_D(\alpha) = \det(T_\alpha)Q_D(\alpha)^2$.*

Proof. Substituting (2) and (3) into (1), we obtain

$$P_D(\alpha) = s\det(T_\alpha) + \sum_{ij\in R_D} \alpha_{ij}^2\det(T_\alpha)t_{ij}^2 + \sum_{\{ij,k\ell\}\in\mathcal{C}(G)} \alpha_{ij}^2\alpha_{k\ell}^2\det(T_\alpha)\big(t_{ij}t_{k\ell} + t_{ik}t_{j\ell} + t_{i\ell}t_{jk}\big)^2.$$

Taking $\det(T_\alpha)$ common gives

$$P_D(\alpha) = \det(T_\alpha)\left(s + \sum_{ij\in R_D} (\alpha_{ij}t_{ij})^2 + \sum_{\{ij,k\ell\}\in\mathcal{C}(G)} \big(\alpha_{ij}\alpha_{k\ell}(t_{ij}t_{k\ell} + t_{ik}t_{j\ell} + t_{i\ell}t_{jk})\big)^2\right).$$

Since squaring distributes over sums in characteristic two and $s^2 = s$, the expression in parentheses is exactly $Q_D(\alpha)^2$ by (4). Therefore, $P_D(\alpha) = \det(T_\alpha) Q_D(\alpha)^2$. □

Lemma 4.2. *If T_α is invertible, then $P_D(\alpha)$ can be evaluated using $O(|V(G)|^3)$ field operations.*

Proof. The determinant and inverse of T_α can be computed by Gaussian elimination using $O(|V(G)|^3)$ field operations. By Lemma 4.1, it is enough to show that $Q_D(\alpha)$ can be computed within the same bound. The constant term and the sum over R_D take only $O(|V(G)|^2)$ operations. The only nontrivial part is the sum over crossing pairs,

$$S = \sum_{\{ij,k\ell\} \in \mathcal{C}(G)} \alpha_{ij}\alpha_{k\ell}\big(t_{ij}t_{k\ell} + t_{ik}t_{j\ell} + t_{i\ell}t_{jk}\big).$$

Every pair of crossing edges can be written uniquely as $ij, k\ell$, where $i < k < j < \ell$. A direct computation of S would therefore use four nested loops, one for each of i, k, j, ℓ, and would take $O(|V(G)|^4)$ operations. We show how to avoid the loop over i. Using the order $i < k < j < \ell$, we can rewrite S as

$$S = \sum_{j} \sum_{k<j} \sum_{\ell>j} \sum_{i<k} \alpha_{ij}\alpha_{k\ell}\big(t_{ij}t_{k\ell} + t_{ik}t_{j\ell} + t_{i\ell}t_{jk}\big).$$

For fixed j, k, ℓ, the factors which do not depend on i can be taken outside the innermost sum. Hence

$$S = \sum_{j} \sum_{k<j} \sum_{\ell>j} \alpha_{k\ell} \left(t_{k\ell} \sum_{i<k} \alpha_{ij} t_{ij} + t_{j\ell} \sum_{i<k} \alpha_{ij} t_{ik} + t_{jk} \sum_{i<k} \alpha_{ij} t_{i\ell} \right).$$

For fixed j, the inner sums can be updated efficiently as k increases. Set $S = 0$. For each j, perform the following steps.

1. Set $h_v \leftarrow 0$ for every vertex v.
2. Process $k < j$ in increasing order. For each such k:
   (a) For each $\ell > j$, set
   $$S \leftarrow S + \alpha_{k\ell}\big(t_{k\ell}h_j + t_{j\ell}h_k + t_{jk}h_\ell\big).$$
   (b) After all $\ell > j$ have been processed, set $h_v \leftarrow h_v + \alpha_{kj}t_{kv}$ for every vertex v.

We now show that the above procedure computes exactly S. Before processing a given k, the value h_v is

$$h_v = \sum_{i<k} \alpha_{ij} t_{iv}.$$

Thus h_j, h_k, and h_ℓ are exactly the three inner sums appearing in the expression for S. Therefore the update of S adds the contribution of all choices $i < k$ for the fixed j, k, ℓ. After processing k, the update $h_v \leftarrow h_v + \alpha_{kj}t_{kv}$ adds the new term corresponding to $i = k$. Hence the same identity for h_v holds for the next value of k. So, as j, k, and ℓ run through all possible choices, we add exactly all terms with $i < k < j < \ell$. These are exactly the terms in the sum defining S. Hence the procedure computes S.

We now count the number of field operations. For each fixed pair j, k, the loop over $\ell > j$ takes $O(|V(G)|)$ field operations. Updating h_v for all vertices v also takes $O(|V(G)|)$ field operations. Since there are $O(|V(G)|^2)$ choices of j, k, the total number of field operations is

$$O(|V(G)|^2) \cdot O(|V(G)|) = O(|V(G)|^3).$$

Thus S, and hence $Q_D(\alpha)$, can be computed using $O(|V(G)|^3)$ field operations. By Lemma 4.1, the same bound holds for evaluating $P_D(\alpha)$. □

Corollary 4.3. *Let G be a graph with a perfect matching, and let D be an orientation of G. There is a randomized test for whether D is Pfaffian using $O(|V(G)|^3)$ field operations per trial. The test always accepts if D is Pfaffian and rejects with probability at least $1/2$ otherwise.*

Proof. Let $|V(G)| = 2n$, and let $q = 2^t$ be the smallest power of two satisfying $q \geq 8n$. Choose $\alpha \in \mathbb{F}_q^{E(G)}$ uniformly at random. We use the following test.

1. If T_α is singular, accept.
2. Otherwise, compute $Q_D(\alpha)$. Reject if $Q_D(\alpha) \neq 0$, and accept otherwise.

If D is Pfaffian, then $P_D(\alpha) = 0$ for every α. If $\det(T_\alpha) = 0$, the test accepts. Otherwise, T_α is invertible, and by Lemma 4.1,

$$P_D(\alpha) = \det(T_\alpha) Q_D(\alpha)^2 = 0.$$

Since $\det(T_\alpha) \neq 0$, we have $Q_D(\alpha) = 0$, so the test also accepts in this case. Hence the test always accepts when D is Pfaffian.

Now suppose that D is non-Pfaffian. Then P_D is a nonzero polynomial of degree $2n$. Since G has a perfect matching, $\det(T)$ is also a nonzero polynomial of degree $2n$. The test can accept only if $P_D(\alpha) = 0$ or $\det(T_\alpha) = 0$. By the Schwartz–Zippel lemma and the union bound,

$$\Pr\big[P_D(\alpha) = 0 \text{ or } \det(T_\alpha) = 0\big] \leq \Pr[P_D(\alpha) = 0] + \Pr[\det(T_\alpha) = 0] \leq \frac{2n}{q} + \frac{2n}{q} = \frac{4n}{q} \leq \frac{1}{2}.$$

Thus, with probability at least $1/2$, T_α is invertible and $P_D(\alpha) \neq 0$. By Lemma 4.1, this implies $Q_D(\alpha) \neq 0$, and the test detects that D is non-Pfaffian.

Testing whether T_α is singular takes $O(|V(G)|^3)$ field operations. If it is invertible, its inverse and $Q_D(\alpha)$ can also be computed within the same bound by the proof of Lemma 4.2. Hence each trial uses $O(|V(G)|^3)$ field operations. □

5 Concluding remarks

We have shown that Pfaffian graph recognition belongs to coRP. For a prescribed orientation, the randomized test can moreover be implemented using $O(|V(G)|^3)$ field operations. The main remaining question is whether Pfaffian graph recognition can be solved in deterministic polynomial time. A related question is whether the polynomial viewpoint can be used directly at the level of the graph, without first passing through the orientation constructed by Vazirani and Yannakakis. Finally, since coRP $\subseteq$ coNP, a coNP-completeness result for Pfaffian graph recognition would imply coNP = coRP and hence NP = RP. Thus, unless NP = RP, Pfaffian graph recognition is not coNP-complete.

AI Disclosure

We obtained the mathematical results independently. ChatGPT (OpenAI) was used for minor language and grammar editing of the manuscript. We take full responsibility for the content.

References

[1] Arthur Cayley. Sur les déterminants gauches. (suite du mémoire t. xxxii. p. 119). *Journal für die reine und angewandte Mathematik*, 38:93–96, 1849. doi:10.1515/crll.1849.38.93.

[2] Marcelo H. de Carvalho, Cláudio L. Lucchesi, and U. S. R. Murty. A generalization of Little's theorem on Pfaffian orientations. *Journal of Combinatorial Theory, Series B*, 102(6):1241–1266, 2012. doi:10.1016/j.jctb.2012.07.004.

[3] Jack Edmonds. Paths, trees, and flowers. *Canadian Journal of Mathematics*, 17:449–467, 1965. doi:10.4153/CJM-1965-045-4.

[4] Ilse Fischer and Charles H. C. Little. A characterisation of Pfaffian near bipartite graphs. *Journal of Combinatorial Theory, Series B*, 82(2):175–222, 2001. doi:10.1006/jctb.2000.2025.

[5] Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge, 2nd edition, 2013. doi:10.1017/CBO9781139020411.

[6] P. W. Kasteleyn. Graph theory and crystal physics. In Frank Harary, editor, *Graph Theory and Theoretical Physics*, pages 43–110. Academic Press, London, 1967.

[7] C. H. C. Little. A characterization of convertible $(0, 1)$-matrices. *Journal of Combinatorial Theory, Series B*, 18(3):187–208, 1975. doi:10.1016/0095-8956(75)90048-9.

[8] László Lovász and Michael D. Plummer. *Matching Theory*, volume 121 of *North-Holland Mathematics Studies*. North-Holland, Amsterdam, 1986.

[9] Cláudio L. Lucchesi and U. S. R. Murty. *Perfect Matchings: A Theory of Matching Covered Graphs*, volume 31 of *Algorithms and Computation in Mathematics*. Springer, Cham, 2024. doi:10.1007/978-3-031-47504-7.

[10] William McCuaig. Pólya's permanent problem. *The Electronic Journal of Combinatorics*, 11(1):R79, 2004. doi:10.37236/1832.

[11] Alberto Alexandre Assis Miranda and Cláudio Leonardo Lucchesi. Recognizing near-bipartite Pfaffian graphs in polynomial time. *Discrete Applied Mathematics*, 158(12):1275–1278, 2010. doi:10.1016/j.dam.2009.01.008.

[12] Serguei Norine. Pfaffian graphs, T-joins and crossing numbers. *Combinatorica*, 28(1):89–98, 2008. doi:10.1007/s00493-008-2150-2.

[13] Serguei Norine and Robin Thomas. Minimally non-Pfaffian graphs. *Journal of Combinatorial Theory, Series B*, 98(5):1038–1055, 2008. doi:10.1016/j.jctb.2007.12.005.

[14] Neil Robertson, Paul D. Seymour, and Robin Thomas. Permanents, Pfaffian orientations, and even directed circuits. *Annals of Mathematics*, 150(3):929–975, 1999. doi:10.2307/121059.

[15] Jacob T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *Journal of the ACM*, 27(4):701–717, 1980. doi:10.1145/322217.322225.

[16] Robin Thomas. A survey of Pfaffian orientations of graphs. In Marta Sanz-Solé, Javier Soria, Juan Luis Varona, and Joan Verdera, editors, *Proceedings of the International Congress of Mathematicians, Madrid 2006, Volume III*, pages 963–984, Zürich, 2006. European Mathematical Society. doi:10.4171/022-3/47.

[17] Leslie G. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8(2):189–201, 1979. doi:10.1016/0304-3975(79)90044-6.

[18] Vijay V. Vazirani and Mihalis Yannakakis. Pfaffian orientations, 0–1 permanents, and even cycles in directed graphs. *Discrete Applied Mathematics*, 25(1–2):179–190, 1989. doi:10.1016/0166-218X(89)90053-X.

[19] Richard E. Zippel. Probabilistic algorithms for sparse polynomials. In Edward W. Ng, editor, *Symbolic and Algebraic Computation: EUROSAM '79*, volume 72 of *Lecture Notes in Computer Science*, pages 216–226, Berlin, Heidelberg, 1979. Springer. doi:10.1007/3-540-09519-5_73.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il