Graphs and Combinatorics (2023) 39:80
https://doi.org/10.1007/s00373-023-02671-7

ORIGINAL PAPER

Undirected Determinant and Its Complexity

Diana Dziewa-Dawidczyk[1] · Adam J. Przeździecki[1]

Received: 8 November 2021 / Revised: 3 June 2023 / Accepted: 4 June 2023 /
Published online: 6 July 2023

Abstract

We view the determinant and permanent as functions on directed weighted graphs and introduce their analogues for undirected graphs. We prove that computing undirected determinants as well as permanents for planar graphs whose vertices have degree at most 4 is #P-complete. In the case of planar graphs whose vertices have degree at most 3, the computation of the undirected determinant remains #P-complete while computing the permanent can be reduced to the FKT algorithm, and therefore can be done in polynomial time. Computing the undirected permanent is a Holant problem and its complexity can be deduced from the existing literature. It is mentioned in the paper as a natural context but no new results in this direction are obtained. The concept of undirected determinant is new. Its introduction is motivated by the formal resemblance to the directed determinant, a property that may inspire generalizations of some of the many algorithms which compute the latter. For a sizable class of planar 3-regular graphs, we are able to compute the undirected determinant in polynomial time.

Keywords Computational complexity · Enumerative combinatorics · Planar graphs · Determinant · Permanent · Pfaffian orientation

Mathematics Subject Classification 05A15 · 05C10 · 05C70

1 Introduction

The most elegant definition of the *determinant* of an $n \times n$ matrix $A = [a_{ij}]$ is obtained when we view A as the adjacency matrix of some weighted directed graph on n vertices, denoted by the same letter A, where the weight of an edge ij is a_{ij}.

✉ Adam J. Przeździecki
adam_przezdziecki@sggw.edu.pl

Diana Dziewa-Dawidczyk
diana_dziewa_dawidczyk@sggw.edu.pl

1 Warsaw University of Life Sciences-SGGW, Warsaw, Poland

A *cycle cover* c of A is a subgraph which contains all the vertices of A and every vertex of c has both in-degree and out-degree 1. Let $cc(A)$ denote the set of all cycle covers of A. Then the determinant of A is given by

$$\det A = (-1)^n \sum_{c \in cc(A)} (-1)^{|c|} w(c), \tag{1.1}$$

and similarly the *permanent* of A equals

$$\operatorname{perm} A = \sum_{c \in cc(A)} w(c), \tag{1.2}$$

where $|c|$ denotes the number of connected components of c, and $w(c)$ is the weight of c, that is, the product of the weights of all the edges in c.

We see that the formulas above also make sense when A is an undirected graph. Thus we define the *undirected determinant* and the *undirected permanent* using formulas (2.1) and (2.2), respectively, applied to an undirected graph A. For undirected graphs, cycle covers are also called 2-*factors*, thus the undirected permanent counts the weighted 2-factors of a graph and the undirected determinant counts them with the sign depending on the parity of the number of components.

We prove in this paper that the problem of computing undirected determinants as well as permanents for planar graphs whose vertices have degree at most 4 is #P-complete (Theorem 3.1). In the case of planar graphs whose vertices have degree at most 3, the computation of the undirected permanent can be reduced (see Sect. 4) to the Fisher–Kasteleyn–Temperley (FKT) algorithm [5, 6, 10], and therefore can be done in polynomial time. This contrasts with our next result, Theorem 5.1, which establishes the #P-completeness of computing the undirected determinant of 3-regular planar graphs. However, for a sizable class of planar 3-regular graphs (actually, we barely peek beyond the bipartite graphs) we are able to compute the undirected determinant in polynomial time (Theorem 6.18). In Appendix we extend the proof of #P-completeness of computing the undirected determinant to the cubic planar graphs whose edges all have weight 1.

Computing the undirected permanent is an instance of a symmetric Holant problem and its complexity can be deduced from the existing literature (see Cai, Fu, Guo and Williams [1]). Nevertheless, we prove the results for the permanent alongside those for the determinant, because little extra work allows us to emphasize the strain between the polynomial and the #P-complete.

The concept of undirected determinant is new. The alterations of the definition of the determinant have a long history. Computing the non-commutative generalization of determinant, called the Cayley determinant, is #P-complete even over the ring of 2×2 matrices (see [3]). Computing the Fermionants and the immanants also tends to be #P-complete (see [4, 8]). It would seem that almost any modification of the definition of determinant leads to a polynomial whose computability is #P-complete. We take a different approach: to draw inspiration from an algorithm rather than from the definition. Although the final definition of the undirected determinant looks like a very natural analogue of the usual determinant, it was actually inspired by the way the closed walks cancel in the Mahajan-Vinay algorithm [7, 9]. The M-V algorithm exploits the fact that sums of closed walks in directed graphs are easily computable and one can arrange a product of such sums where all the intersecting cycles cancel, leaving precisely the cycle covers. A naive undirected version of M-V would use sums of closed non-backtracking walks, which are also easily computable. In the analogous products of these sums, at least in the case of cubic graphs, many intersecting cycles cancel, but not all. In this paper we go some distance toward computing u-det for some graphs in polynomial time and hope that more is possible.

Another source of our motivation is the intuition that the main reason that makes these alterations #P-complete is their deprivation of some of the very special properties enjoyed by the determinant. More specifically, we suspect that the key limitations of the FKT algorithm stem from its direct dependence on the determinant. Linear properties of the determinant cause the FKT algorithm to count only those structures which allow some kind of "interpolation". To illustrate what we have in mind, let us recall that Valiant introduced matchgates [12, Section 2.2], [13, Section 4]. Then, the characterization of the possible signatures of planar matchgates has been completed by Cai and Gorenstein [2] in terms of Matchgate Identities. These identities imply that for any matchgate G, if its signature Γ_G is non-zero on two length-k bitstrings $\alpha, \beta \in \{0, 1\}^k$ then there exists a sequence

$$\alpha = \alpha_0, \alpha_1, \ldots, \alpha_s = \beta,$$

in $\{0, 1\}^k$ such that α_{i-1} differs from α_i at exactly two places for $i = 1, 2, \ldots, s$ and $\Gamma_G^{\alpha_i} \neq 0$ for $i = 0, 1, \ldots, s$. This contrasts with the behavior of Boolean formulas: the information that two assignments $(x_1, x_2, \ldots, x_n)$ and $(y_1, y_2, \ldots, y_n)$ satisfy a formula ϕ does not imply that any other assignment does.

The reach of the FKT algorithm has been extended by holographic reductions introduced by Valiant [14, 15]. These, at least in large part, are a way to incorporate the linear properties of the determinant into the formulation of combinatorial problems.

The undirected determinant seems promising since on the one hand it exhibits essentially none of the linear properties of the determinant, which we like, but which make it difficult for many combinatorial counting problems to be evaluated with the usual determinant. Computing u-det is #P-complete even on cubic planar graphs. On the other hand, u-det does not stray very far from some of the algorithms for det, which may raise hopes of adapting them to compute u-det.

2 Notation and Preliminaries

Most of our terminology is standard and follows for example Thomas [11]. All graphs considered are finite, planar and undirected, although in Sect. 6 we consider undirected graphs whose edges are equipped with an orientation. The main graphs constructed in this paper have no loops or multiple edges. Graphs with loops or multiple edges may appear only in side comments or reasonings. Each case when they actually occur is explicitly mentioned in the text. In each such case either we explicitly show how to modify our graphs to remove loops and multiple edges, or these are used only as parts of a proof.

The symbols $V(G)$ and $E(G)$ denote the sets of vertices and edges of a graph G. We write $v \in G$ or $e \in G$ instead of $v \in V(G)$ or $e \in E(G)$ when no confusion can arise. Elements of $E(G)$ are denoted $\{v_1, v_2\}$, or $v_1 v_2$ if we want to indicate that the edge is oriented from v_1 to v_2. The graph G is *weighted*, which means that it is equipped with a weight function $w : E(G) \to F$. The reader may view the F as the rational numbers; we are going to use only the weights 1, -1 and, in Sect. 5, $-\frac{1}{2}$. However, in our constructions, F may be any field of characteristic other than 2.

A *cycle cover* of G, denoted $cc(G)$, is a spanning subgraph whose vertices all have degree 2. In the literature, cycle covers are also called vertex cycle covers or 2-factors. The weight of a cycle cover $c \in cc(G)$ is the product of the weights of its edges: $w(c) = \prod_{e \in c} w(e)$.

The *undirected determinant* and the *undirected permanent* are defined as

$$\text{u-det}\, G = (-1)^n \sum_{c \in cc(G)} (-1)^{|c|} w(c), \tag{2.1}$$

and

$$\text{u-perm}\, G = \sum_{c \in cc(G)} w(c), \tag{2.2}$$

where n is the number of vertices in G and $|c|$ is the number of cycles in c. These definitions make sense also in the case of graphs with loops and multiple edges.

The remainder of this section is devoted to introducing gadgets, convenient components that facilitate the construction of larger graphs. Permanental gadgets are little different than those used in matching theory; their properties are independent of the planarity of the external part of the graph. On the other hand, determinantal gadgets need a little extra structure and require the ambient graph to be planar.

Definition 2.3 A *gadget* is a graph G_0 equipped with a set $ext(G_0)$ of *external edges* such that each $e \in ext(G_0)$ is adjacent to exactly one vertex in G_0. We consider only those gadgets G_0 which are planar and whose external edges stretch towards the unbounded face. When referring to gadgets, we use the word *cycle* to mean either an actual cycle or a path connecting two external edges. Each cycle cover $c \in cc(G_0)$ of a gadget G_0 determines a, possibly empty, subset $S \subseteq ext(G_0)$ consisting of those edges in $ext(G_0)$ which belong to c.

Definition 2.4 A *determinantal gadget* is a gadget augmented by a choice of pairings of elements of $S \subseteq ext(G_0)$. For every subset S of even size we choose a pairing of its elements (= decomposition into pairs), denoted by P_S. Each of these pairings must be realized by disjoint paths connecting an element of S to another element of S. To increase the flexibility of the reasonings, these paths may be any disjoint simple curves in the plane $\mathbb{R}^2$ in which the gadget is embedded. These paths avoid the outer face of the gadget, however, we do not require them to be graph-theoretic paths.

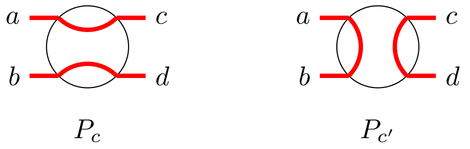

Fig. 1 Modification of a pairing of external edges inside a gadget

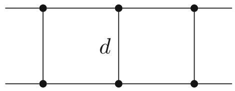

Fig. 2 The skew crossover gadget (weight 1 edges are unlabeled)

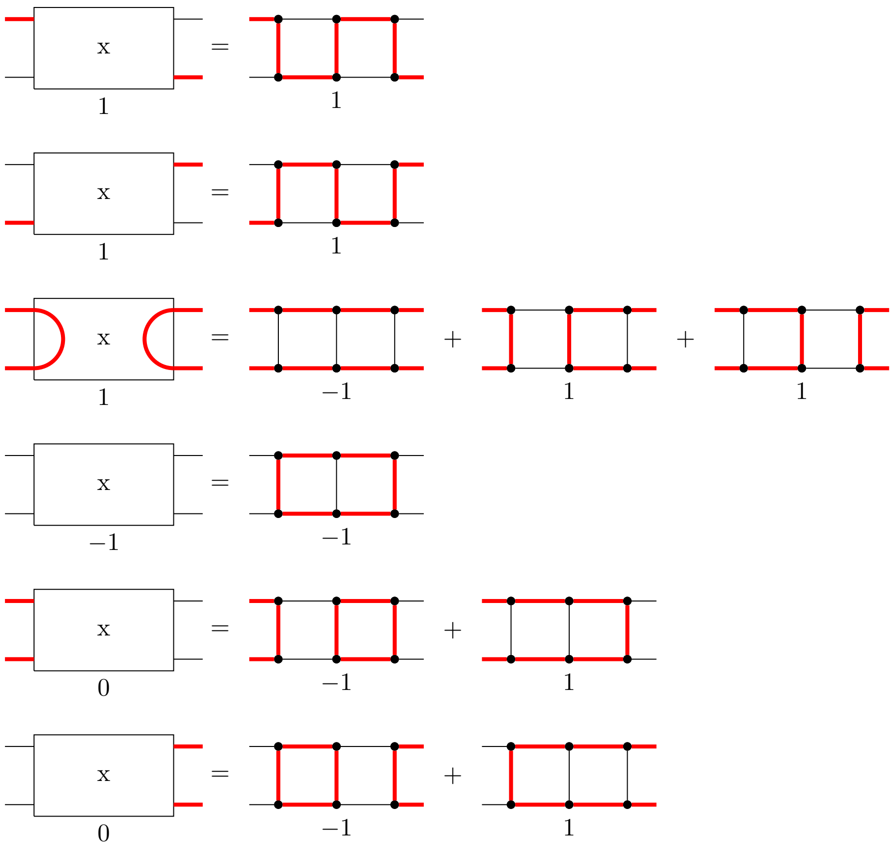

Fig. 3 The signature of the skew crossover gadget in the case of determinant; the permanental signature has the opposite sign

The sets S of size 2 admit the unique pairing so that there is no need to indicate it. The pairing P_S is clearly marked in the gadget shown in Fig. 3. The remaining gadgets constructed in the paper are of two types: those which synchronize edges (Figs. 5, 7 and 8), they have a natural choice of P_S; and those where the choice of P_S is irrelevant since the signature, constructed in Definition 2.6 below, is 0 whenever $|S| > 2$ (all gadgets in Sect. 5).

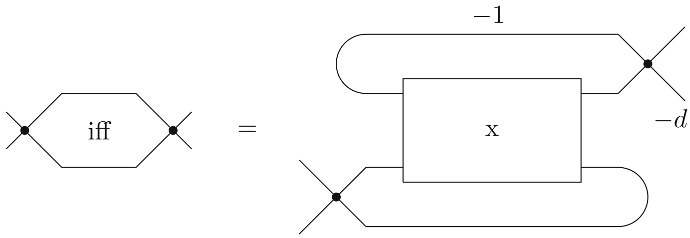

Fig. 4 The iff gadget and its symbolic notation

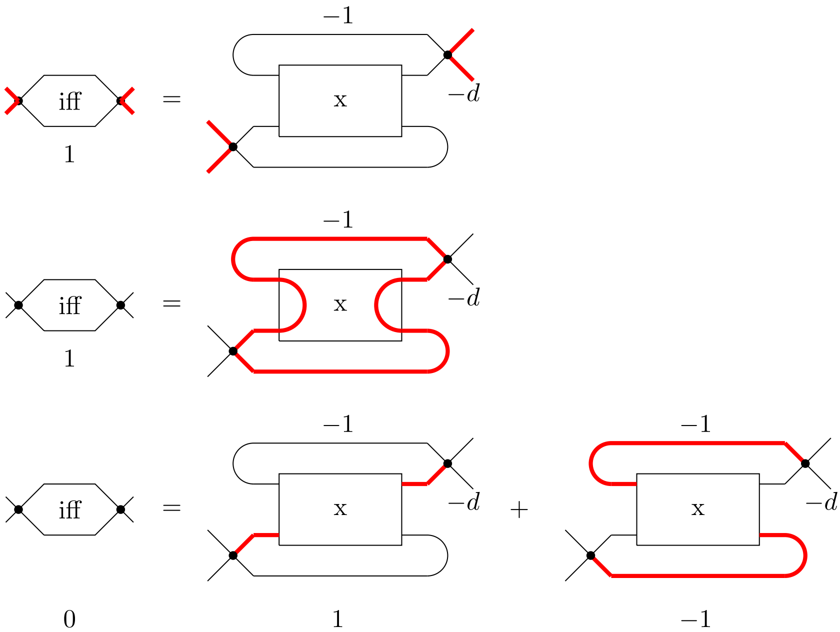

Fig. 5 Nonzero signatures of the iff gadget and cancellation of left to right paths. When all the external edges belong to a cycle cover of the ambient graph, we determine the parity of the number of non-inner cycles by connecting the left two and the right two edges

If $|S| > 2$, a pairing, as above, may be modified as shown in Fig. 1. Each modification involves two pairs and is of the form $\{\{a, c\}, \{b, d\}\} \mapsto \{\{a, b\}, \{c, d\}\}$ for some $a, b, c, d \in S$.

Lemma 2.5 *If P_1 and P_2 are two pairings, as defined above, of the same subset $S \subseteq ext(G_0)$ then they can be connected by a sequence of modifications shown in Fig. 1.*

The parity of the number of modifications in a sequence which connects P_1 and P_2 does not depend on the sequence.

Proof Enumerate the elements of $S = \{e_1, e_2, \ldots, e_{2k}\}$ in the order they appear on the boundary of the gadget. Let P_0 be the pairing obtained by connecting every e_{2i-1}

to e_{2i} for $i = 1, 2, \dots, k$. Clearly P_0 can be realized by disjoint paths drawn on the plane inside the gadget.

In order to prove the first part of the statement it is enough show how to connect P_1 to P_0. We proceed by induction. If two consecutive elements of S, e_{2i-1} and e_{2i}, are already connected in P_1, we do nothing. Otherwise, they are ends of two paths which connect $\{e_{2i-1}, x\}$ and $\{e_{2i}, y\}$ for some $x, y \in S$. Since consecutive elements cannot be separated by a path in P_1, we may apply the modification shown in Fig. 1 for $a = e_{2i-1}$, $b = e_{2i}$, $c = x$ and $d = y$. Eventually we obtain P_0.

In order to prove the second part of the statement we connect the elements of S by disjoint paths drawn on the plane outside the gadget. We use the same outer paths for P_1 and for P_2. The outer paths together with P_1 form n_1 disjoint circles on the plane and similarly n_2 circles when completed with paths in P_2. Every modification in Fig. 1 changes the parity of the number of circles, hence the parity of the number of modifications is the same as the parity of $n_2 - n_1$. □

At the end of the proof above, we claim that every modification changes the parity of the number of circles. Note that this is valid for planar graphs only.

A cycle cover $c \in cc(G_0)$ consists of cycles and paths with ends in $ext(G_0)$. These paths determine a pairing, denoted by P_c, of the set $S = c \cap ext(G_0)$. We use the choice of P_S to define a function

$$\tau : cc(G_0) \to \{0, 1\}$$

as $\tau(c) = 0$ if P_c can be transformed to P_S by an even number of modifications and $\tau(c) = 1$ otherwise. Lemma 2.5 implies that τ counts the parity of the number of modifications connecting P_c and P_S.

Definition 2.6 A *signature* of a gadget G_0 is a function

$$\text{signature}_* : \mathcal{P}(ext(G_0)) \to F$$

where $\mathcal{P}(ext(G_0))$ denotes the set of all subsets of $ext(G_0)$.

The *determinantal signature* of a gadget G_0 equipped with a function $\tau : cc(G_0) \to \{0, 1\}$ is defined as

$$\text{signature}_d(S) = (-1)^{n_0} \sum_{\substack{c \in cc(G_0) \\ c \cap ext(G_0) = S}} (-1)^{|c| + \tau(c)} w(c),$$

while the *permanental signature* is

$$\text{signature}_p(S) = \sum_{\substack{c \in cc(G_0) \\ c \cap ext(G_0) = S}} w(c),$$

where n_0 is the number of vertices of G_0 and $|c|$ is the number of cycles in c. In particular, when $|S|$ is odd, the sums are empty and the signature at such S is 0.

The signature allows us to compute the undirected determinant (or permanent) of a planar graph which contains G_0 as a subgraph without analyzing the, possibly complicated, inner structure of G_0. This is explained by the following.

Lemma 2.7 *Let $G_0 \subseteq G$ be a gadget in a planar undirected weighted graph. Let $G_\infty = G \setminus G_0$ and define the cycle covers $cc(G_\infty)$ analogously to $cc(G_0)$. Let n_∞ be the number of vertices in G_∞.*

Then

$$\text{u-det}\, G = (-1)^{n_\infty} \sum_{S \subseteq ext(G_0)} \text{signature}_d(S) \sum_{\substack{c \in cc(G_\infty) \\ c \cap ext(G_\infty) = S}} (-1)^{|c|} w(c) \tag{2.8}$$

and

$$\text{u-perm}\, G = \sum_{S \subseteq ext(G_0)} \text{signature}_p(S) \sum_{\substack{c \in cc(G_\infty) \\ c \cap ext(G_\infty) = S}} w(c)$$

The formulas above involve, respectively, the determinantal and the permanental signatures.

The symbol $|c|$ denotes the number of cycles in $c \in cc(G_\infty)$ obtained when c is completed with the paths in the pairing P_S was introduced in Definition 2.4.

Proof The proof is straightforward. The sums indexed over $c \in ext(G)$ in formulas (2.1) and (2.2) are grouped according to $S = c \cap ext(G_0)$. Then for every S, the weights of the edges that belong to G_0 are factored out into the $\text{signature}_*(S)$.

In the case of u-det we need to verify the signs. The number $|c|$ in formula (2.1) splits into the sum of the number of cycles that are entirely inside of the gadget (those are counted by the signature) and the number of remaining cycles, which are modified inside the gadget by P_S and then counted by formula (2.8). The cases when this modification changes parity are corrected by $\tau(c)$ which is used to define the signature (see Definition 2.6). □

By abuse of terminology, we often write “signature” when we mean the value of the signature at a specific subset. When we list the values of the signature we may omit subsets at which the signature is zero.

3 Computing the Undirected Determinant and Permanent of Planar Graphs of Maximum Degree 4 is #P-complete

In this section, for a given Boolean formula $\phi(x_1, x_2, \ldots, x_n)$ in conjunctive normal form with m clauses, where each clause is limited to three literals, we construct undirected weighted planar graphs A_ϕ and B_ϕ, of maximum degree 4, such that

- the number of assignments satisfying ϕ,
- $(-1)^m$ times the undirected determinant of A_ϕ, and

Fig. 6 A version of the iff gadget when multiple edges are allowed

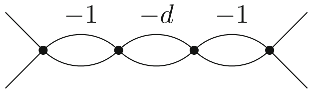

– the undirected permanent of B_ϕ

are all equal. The graphs A_ϕ and B_ϕ differ only by signs of weights of some edges, thus it is convenient to describe them in parallel. The size of each of these graphs is $O(m^2)$.

We construct these graphs by means of the following gadgets:

1. The skew crossover gadget.
2. The iff gadget: it synchronizes two edges which belong to the boundary of a common face of a planar graph. Connecting the two edges with an iff gadget results in a new graph whose permanent (respectively determinant) counts the weights of only those cycle covers of the original graph which contain either both or none of the edges connected by the gadget.
3. The extended iff gadget: it synchronizes any two edges of a planar graph.
4. The variable setting gadget: it encodes the variables of the formula ϕ.
5. The clause gadget: it encodes the clauses of the formula ϕ.

(1) *The skew crossover gadget.* We start with a skew crossover gadget, shown in Fig. 2. We follow the convention that all unlabeled edges have weight 1. The gadget is inspired by the Cai–Gorenstein [2] construction. Its signature is 0 unless the opposite edges, either both or none, belong to a cycle cover. The nonzero signatures are either 1 or -1.

The weight d is defined as $d = 1$ for determinant and $d = -1$ for permanent.

In Fig. 3, we list all the possible ways a cycle cover may meet this gadget, and the ways they add to yield the signature of this gadget, in the permanental case.

In the determinantal case, all values of the signature of this gadget have the opposite sign, but the third one requires an additional comment. This is the only instance where the parity of the number of components of the cycle cover depends on the way it meets the gadget. The arcs inside the box indicate the way the cycles should traverse this gadget, for the sake of counting the parity of the number of cycles in a cover.

(2) *The iff gadget*, shown in Fig. 4, has four external edges and its signature is 1 if either all or none of them belong to a cycle cover. Otherwise its signature is 0. All other cycle covers cancel out. The gadget is used to synchronize two edges which belong to the boundary of a common face of the planar graph.

The iff gadget can be constructed using the skew crossover gadget, as shown in Fig. 4.

Notice the $-d$ label near one of the external edges. It indicates that, when applying the iff gadget, the original weight of this edge has to be multiplied by $-d$. As above, we have $d = 1$ in the case of determinant and $d = -1$ for permanent.

If a component of a cycle cover c passes left to right through the gadget then, by its symmetry, there exists a different cycle cover c' with $w(c') = -w(c)$, thus c and c' cancel out.

The remaining, easy to compute, nonzero signatures of the iff gadget are shown in Fig. 5.

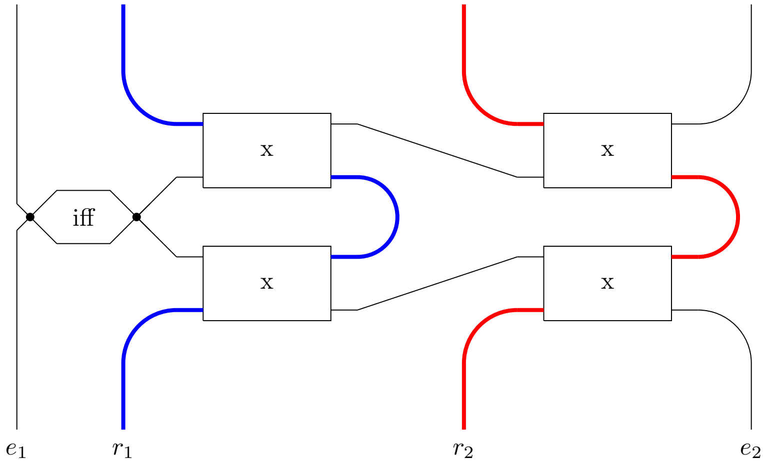

Fig. 7 The extended iff gadget synchronizing two edges e_1 and e_2, across two other edges r_1 and r_2. The colors indicate the way the paths, representing r_1 and r_2, traverse the gadget when both the r_i's belong to the cycle cover but none of the e_i's does (color figure online)

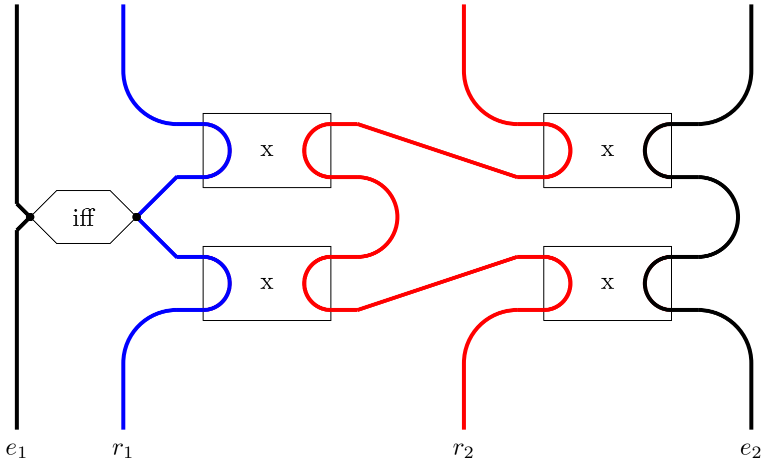

Fig. 8 The same gadget as in Fig. 7. The case when all the r_i's and the e_i's belong to the cycle cover

In the determinantal case, the inner loop introduces the -1 sign. This loop is not seen outside the iff gadget hence, its sign is included in the signature.

In Fig. 4, we give a construction that avoids loops and multiple edges. However, if we do not have to avoid multiple edges, we may use a simpler and more obvious construction of the iff gadget, shown in Fig. 6.

(3) *The extended iff gadget.* We use the skew crossing gadgets and the iff gadget, defined above, to synchronize any two edges e_1 and e_2 in the graph. By the synchronization of the edges e_1 and e_2 we mean that inserting the *extended iff gadget* into the graph causes the determinant (resp. permanent) of the new graph to count precisely those cycle covers of the original graph which contain either both or none of the edges

Fig. 9 A symbolic notation for the gadget shown in **Figs.** 7 and 8

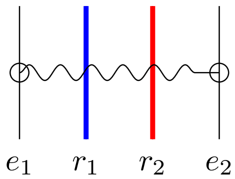

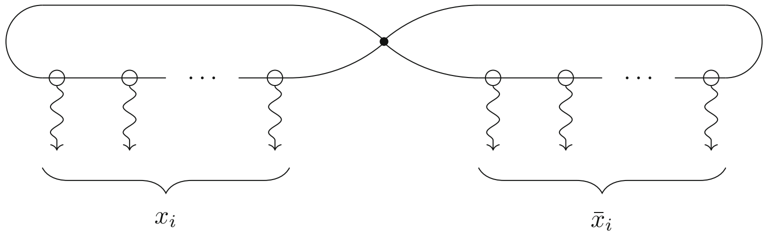

Fig. 10 The variable setting gadget

e_1 and e_2. Otherwise, the gadget has no effect on the remaining cycle covers. In Figs. 7 and 8 we present the construction of the extended iff gadget that goes across two other edges.

The remaining cases, when the number of edges r_i, $i = 1, 2, \ldots, r_n$, the gadget passes through is different from 2, as well as different configurations of which of the r_i's belong or not to the cycle cover, are analogous.

Figure 9 shows the symbolic notation of the gadget.

We see that, any two skew crossing gadgets, that are drawn one above the other in Figs. 7 and 8, have the same signature, either 1 or -1. Therefore the crossing gadgets do not contribute to the values of the signature of the extended iff gadget. Neither does the iff gadget.

It is straightforward to see that those cycle covers of the modified graph that contribute to either the determinant or permanent must contain either both or neither of the edges e_1 and e_1.

It remains to notice that insertion of the extended iff gadget affects neither the number of cycle covers nor the parity of the number of cycles in a cycle cover. The only restriction is that e_1 belongs to a cover if and only if e_2 does.

When neither e_1 nor e_2 is in the cover, a case similar to the one shown in Fig. 7, whether a path representing r_i belongs to the cover or not does not affect other paths.

When both e_1 and e_2 are in the cover, we start with a case similar to the one shown in Fig. 8, where the cycle cover passes through all the r_i's, $i = 1, 2, \ldots, n$. Removal of any path representing an edge r_i from the cycle cover causes its central part to be filled by the next path on the right.

We conclude that whenever the extended iff gadget is placed between any two edges e_1 and e_1 as in Fig. 9, the determinant as well as the permanent of the resulting graph counts exactly those cycle covers of the original graph which contain either both or neither of the e_i's.

(4) *The variable setting gadget.* For every variable x_i in the formula ϕ we construct a gadget as shown in Fig. 10.

Fig. 11 The clause gadget encoding $a \vee b \vee c$

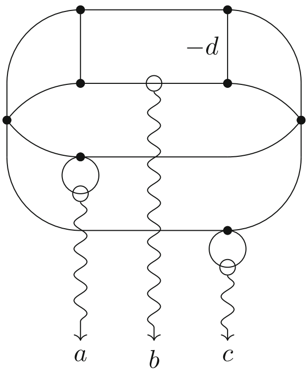

The gadget consists of two loops connected to a single vertex. Clearly every cycle cover contains exactly one of the two loops and the signature of the gadget is always 1 so that addition of the variable setting gadget does not affect the determinant or permanent of the graph. One of these two loops represents the variable x_i, while the other represents its negation $\bar{x}_i$. The wavy lines in Fig. 10 denote the extended iff gadgets which synchronize a loop with those edges in the clause gadgets, shown in Fig. 11, which represent the same x_i or $\bar{x}_i$, respectively. We have one extended iff gadget per one occurrence of x_i or $\bar{x}_i$ in the formula ϕ.

If x_i or $\bar{x}_i$ is not present in ϕ then the graph we obtain has a loop. If this is undesirable we may synchronize this loop with itself, using the iff gadget, which results in a loopless graph.

(5) *The clause gadget.* For every clause of the form $a \vee b \vee c$ in the formula ϕ we construct a gadget as shown in Fig. 11.

Three edges of the gadget represent the three literals $a, b, c \in \{x_1, \bar{x}_1, x_2, \bar{x}_2, \ldots\}$ in the clause. The wavy lines indicate the extended iff gadgets which synchronize these three edges with those loops in the variable setting gadgets which represent the same literal.

All the possible cycle covers of the clause gadget are shown in Fig. 12. We see that the signature is 0 if all the three literals are set to *false*, represented by 0. Otherwise, all the signatures are 1 in the permanental case and -1 in the determinantal case.

The construction of this gadget was chosen so as to make the exposition more straightforward; however, the gadget can be simplified by removing those edges that correspond to literals a and c and connecting the respective extended iff gadgets to $\bar{a}$ and $\bar{c}$.

Theorem 3.1 *Computing the undirected determinant and the undirected permanent of planar graphs whose vertices have degrees* 3 *or* 4 *is #P-complete.*

Proof Consider a Boolean formula $\phi(x_1, x_2, \ldots, x_n) = c_1 \wedge c_2 \wedge \cdots \wedge c_m$, where $c_j = t_{j,1} \vee t_{j,2} \vee t_{j,3}$ with $t_{j,k} \in \{x_1, \bar{x}_1, x_2, \bar{x}_2, \ldots, x_n, \bar{x}_n\}$, where $j = 1, 2, \ldots, m$ and $k = 1, 2, 3$. Here $\bar{x}_i$ represents the negation of x_i.

We construct the planar graphs A_ϕ for determinant and B_ϕ for permanent by taking one clause gadget for every c_i and one variable setting gadget for every x_i. For every

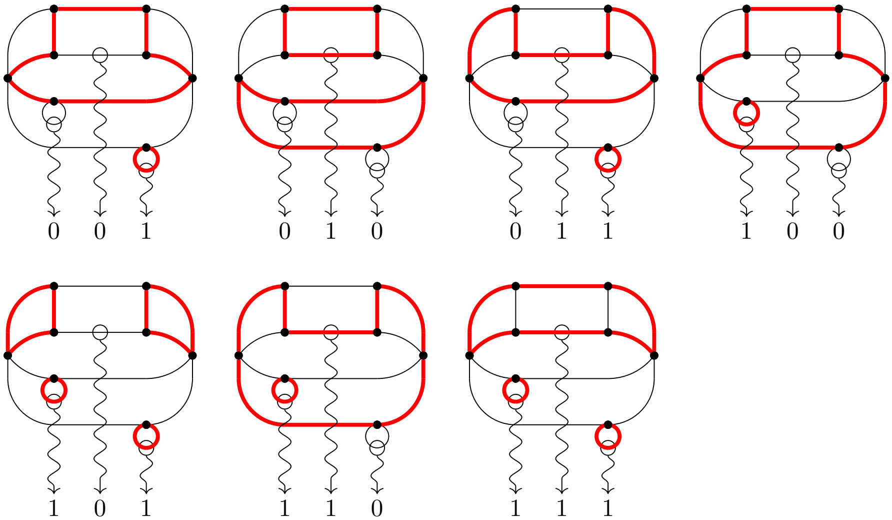

Fig. 12 All the 7 cycle covers of the clause gadget. The labels indicate the corresponding valuations of a, b, c

occurrence of x_i or $\bar{x}_i$ in the formula, we use the extended iff gadget to synchronize a suitable edge of the variable setting gadget with the corresponding edge in the clause gadget.

The properties of the gadgets, described in this section, imply that the number of satisfying valuations of ϕ is equal to $(-1)^m$ times the undirected determinant of A_ϕ and to the undirected permanent of B_ϕ. This proves the theorem. □

As an aside, let us notice that the only weights used in this section are -1, 0 and 1. Therefore the reduction above can be applied to any field F, in which case we obtain the number of satisfying valuations modulo the characteristic of the field. In particular, for $\mathrm{ch}(F) = 2$, computing the undirected determinant is a $\oplus$P-complete problem. Therefore one does not expect it to be polynomially computable even in this case.

4 Computing the Undirected Permanent of Planar Graphs of Maximum Degree 3 Reduces to the FKT Algorithm

For the sake of completeness we present a straightforward observation that computing the undirected permanent of cubic planar graphs reduces to the FKT algorithm. Consequently, over a field of characteristic 2, also computing the undirected determinant reduces to FKT.

Let G be a weighted undirected planar graph of maximum degree 3. If G contains a vertex of degree 0 or 1 then clearly

$$\text{u-perm}\, G = 0.$$

Suppose that all the vertices of G have degree 2 or 3 and all the edges have nonzero weight. We define G_{inv} as the subgraph of G induced by all vertices of degree 3, and the weights of G_{inv} are inverted, that is, $w_{G_{inv}}(e) = 1/w_G(e)$, where e denotes an edge and $w_G(e)$ denotes its weight in the graph G. The complement of a cycle cover of G is a perfect matching in G_{inv}. In fact the complement, in the set of edges of G, establishes a one-to-one correspondence between the perfect matchings in G_{inv} and the cycle covers of G. The product of the weight of a cycle cover and the corresponding perfect matching is always equal to the product of the weights of all the edges of G, denoted by p. Since the weights in G_{inv} are inverted, we see that

$$\text{u-perm}\, G = p \cdot \text{PerfMatch}\, G_{inv},$$

where PerfMatch G_{inv} is the sum of the weights of all the perfect matchings of G_{inv}. Since G_{inv} is planar, the latter sum is computed in polynomial time by the FKT algorithm.

5 Computing the Undirected Determinant of Cubic Planar Graphs is #P-Complete

Unlike the undirected permanent, it turns out that computing the undirected determinant is #P-complete even in the case of cubic planar graphs. By Theorem 3.1 we already know that computing the undirected determinant is #P-complete for planar graphs whose vertices have degree 3 or 4, hence it is enough to construct a cubic planar gadget whose signature is the same as that of a single vertex of degree 4. This reduction involves multiplying by powers of 2 and therefore requires the field over which the edges are labeled to have characteristic different from 2. In view of Sect. 4 this is unavoidable.

Let us point out that while all the gadgets constructed in Sect. 3 had edges of weight either 1 or -1, here we need to know that 2 is invertible or, at least, it is a non-zero-divisor. This restriction is impossible to avoid since, modulo 2, the undirected determinant coincides with the undirected permanent, and computing the latter reduces to the FKT algorithm, as proved in Sect. 4.

When defining a determinantal gadget one should choose a pairing P_S for every set S of an even number of external edges of the gadget. However, the signatures of the gadgets constructed in this section are null whenever $|S| > 2$. Therefore these signatures are independent of the choice of such pairings and we could conveniently ignore discussion of the pairings.

(1) *The auxiliary gadget* is shown in Fig. 13.

The signatures, up to symmetry, corresponding to the different ways a cycle cover can traverse the auxiliary gadget are shown in Fig. 14. The last (fourth) signature is included only for the sake of completeness, as it is never realized in the subsequent constructions. The value s in the computation of this signature is either 1 or -1, depending on which way of passing through the gadget is chosen to be positive. We do not specify this choice since in any case this signature is 0.

(2) *The null edge gadget* is shown in Fig. 15. This gadget plays the role of an edge

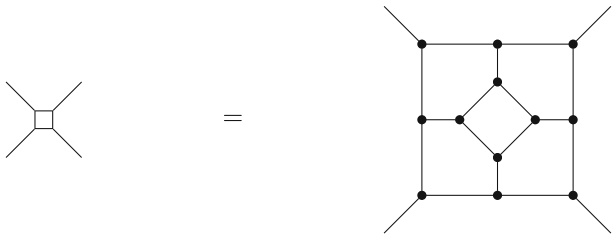

Fig. 13 The auxiliary gadget and its symbolic notation

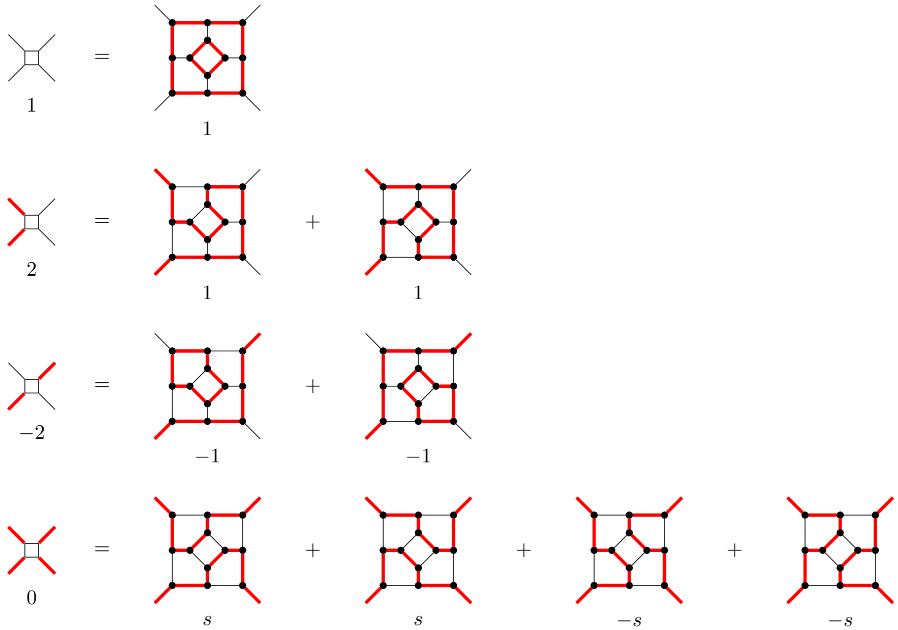

Fig. 14 The signatures, up to symmetry, of the auxiliary gadget

of weight 0, so that the gadget is not necessary but convenient if we want to ensure that all vertices have degree 3 and all weights are invertible. The arguments employed in Sect. 6 are neater when we deal with cubic graphs instead of graphs of degree at most 3.

(3) *The degree* 4 *vertex gadget* is constructed in Fig. 16.

The signatures of the degree 4 vertex gadget are listed, up to symmetry, in Fig. 17.

Theorem 3.1 and the existence of a planar 3-regular gadget whose signature is equal to -4 times the signature of a single vertex of degree 4 implies the following.

Theorem 5.1 *Computing the undirected determinant of cubic planar graphs with weights in the set* $\{-1, -\frac{1}{2}, 1\}$ *is #P-complete.*

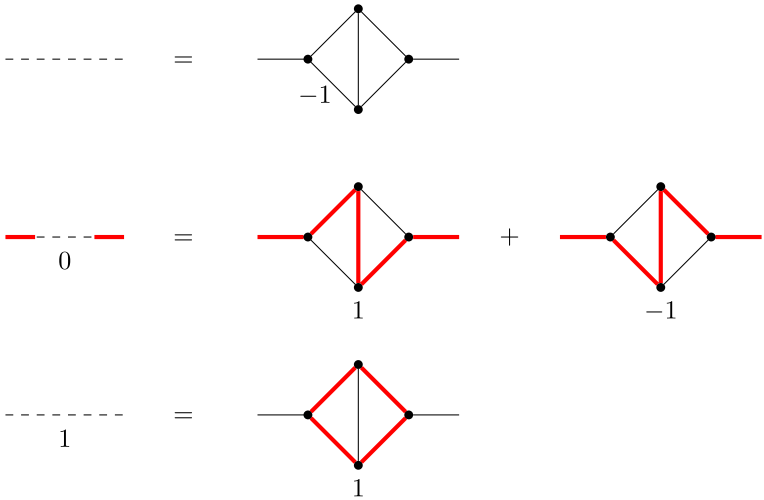

Fig. 15 The null edge gadget and its signature

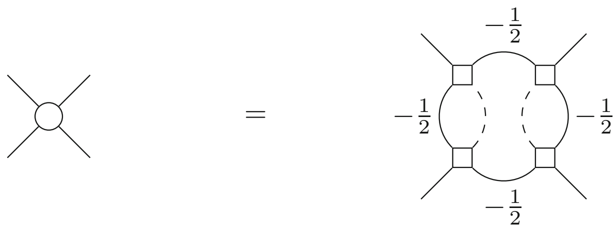

Fig. 16 The degree 4 vertex gadget and its symbolic notation

6 Semi-Pfaffian Orientation and Computing the Undirected Determinant

The purpose of this section is twofold. Firstly, we prove Theorem 6.18 which states that the undirected determinant is polynomially computable for a reasonable class of cubic planar graphs, which includes the bipartite graphs. This, together with Theorem 5.1, introduces another instance of a tension between P and #P. The second aim of this section is the search for tools and ideas which could guide us in our attempts to find polynomially computable analogues of the determinant, hopefully improving the computational strength of the FKT algorithm. In this direction, Definition 6.3 introduces the semi-Pfaffian orientation and Definition 6.8 introduces the tension of an even cycle in a planar graph. Both play the key role in our proof of the polynomial computability for the graphs mentioned above.

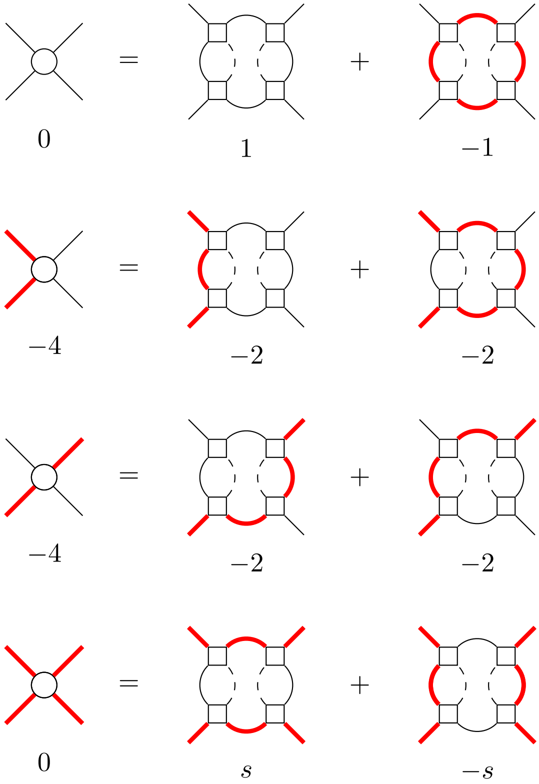

Fig. 17 The signatures, up to symmetry, of the degree 4 vertex gadget

Additional Preliminaries

Let $G = (V, E)$ be a weighted undirected graph. An *orientation* of G is a choice, independently for every edge of G, of a direction from one of its end points to the other; the orientation is not considered a part of the graph structure. An undirected graph G on vertices $V = \{1, 2, \ldots, n\}$, equipped with an orientation, is represented by a skew symmetric matrix $A = [a_{ij}]_{i,j=1,2,\ldots,n}$ where, for every edge of weight e, oriented from i to j, we have $a_{ij} = e$ and $a_{ji} = -e$. We put $a_{ij} = 0$ if there is no edge $\{i, j\}$ in G.

A cycle c in a graph G is *even* if it has even length, and *central* if $G \backslash V(c)$ has a perfect matching. An even cycle c is *oddly oriented* if for either choice of direction of traversal around c, the number of edges of c directed according to the direction of the traversal is odd.

Occasionally, it is convenient to abuse the notation and treat a planar graph as if it was a subset of a plane.

Recall that an orientation of the edges of G is *Pfaffian* if every even central cycle of G is oddly oriented. At the heart of the FKT algorithm we see two theorems:

1. Every planar graph admits a Pfaffian orientation.

2. If A is the skew symmetric adjacency matrix, associated with a graph G with Pfaffian orientation, then, up to sign,

$$\text{Pfaffian } A = \text{PerfMatch } G.$$

The standard definition of Pfaffian (see for example Thomas [11, Section 2]) involves a sign function $\text{sgn}_G : pm(G) \longrightarrow \{-1, 1\}$, defined on the set $pm(G)$ of all perfect matchings of G. It depends on the orientation of G and is given by

$$\text{sgn}_G(a) = \text{sgn} \begin{pmatrix} 1 & 2 & 3 & 4 & \dots & 2n-1 & 2n \\ i_1 & j_1 & i_2 & j_2 & \dots & i_n & j_n \end{pmatrix} \tag{6.1}$$

where sgn is the sign of the permutation, and the edges of the perfect matching $a = \{i_1 j_1, i_2 j_2, \dots, i_n j_n\}$ are listed in such a way that every edge $i_k j_k$ is directed from i_k to j_k.

The Pfaffian of A is defined as

$$\text{Pfaffian } A = \sum_{a \in pm(G)} \text{sgn}_G(a) w(a). \tag{6.2}$$

Let us recall that the Pfaffian is polynomially computable. In fact, the algorithms that compute the determinant tend to translate to algorithms for the Pfaffian.

The Semi-Pfaffian Orientation

Definition 6.3 An orientation of a graph G is *semi-Pfaffian* if a central cycle in G of length $2k$ is oddly oriented if and only if k is odd.

Remark 6.4 Not all planar graphs admit semi-Pfaffian orientation, but those which have at most two faces bounded by an odd number of edges do. This includes the bipartite graphs.

Proof If no faces of the graph have an odd number of edges then the graph is bipartite and we orient the edges from one side of the bipartition to the other and the claim is obvious.

Otherwise, since every edge belongs to precisely two faces, we must have two odd faces (one may be unbounded). Connect the interior of one odd face with the other using a path P in the plane which avoids vertices and crosses each edge at most once.

We split vertices of the graph into two subsets according to the following rule. Both ends of an edge which crosses P belong to the same subset. The ends of an edge disjoint from P belong to distinct subsets. Such a partition exists since P intersects the boundary of the end faces an odd number of times, and the remaining faces an even number of times. Therefore, whenever we travel along the perimeter of any face, we switch the subsets an even number of times, and therefore this partition is well defined.

For those edges which are disjoint from P we choose the orientation from one side of this partition to the other. For those edges which cross P we choose the orientation from one side of the path to the other.

Let c be an even cycle in G. It must cross P an even number of times. We choose a direction of traversal around c and claim that c crosses the path P the same number of times in either direction. If not, we could find two points p_0 and p_1 in P such that c crosses the path in the same direction in both of these points and the interior of the segment $Q \subseteq P$ between p_0 and p_1 is disjoint from c. Let $C \subseteq c$ be the plane path which is a segment of c between p_0 and p_0. We see that $C \cup Q$ is a closed curve which splits the plane into two disjoint components. The cycle c enters C from one component and leaves C towards the other. This is impossible since c is a cycle without self-intersections and is disjoint from Q.

We conclude that c has an even number of edges that cross P; denote this number $2a$. The orientation introduced in the fourth paragraph of the proof is such that a of these edges are directed in either of the two directions of traversal around c. When we contract these $2a$ edges we see that any two vertices which became identified belong to the same side of the partition. We are left with $2b$ edges which were disjoint from P, and see that the partition constructed in the third section induces a bipartition of the contracted cycle. Therefore b of these edges are directed in either of the two directions of traversal around c. This completes the proof. □

Cubic Planar Graphs

From now on, we assume that G is a weighted undirected cubic planar graph with invertible weights of edges. G is equipped with a semi-Pfaffian orientation, and is represented by a skew symmetric matrix A.

In such a graph, the complement $\bar{a}$ of a perfect matching $a \in pm(G)$ in the set of edges $E(G)$ is a cycle cover. Conversely, the complement $\bar{c}$ of a cycle cover $c \in cc(G)$ is a perfect matching. Note that if p is the product of all the weights of the edges in G then for every $c \in cc(G)$ we have

$$w(c)w(\bar{c}) = p. \tag{6.5}$$

Since we are ultimately interested in the undirected determinant, we rewrite formula (6.2) for the Pfaffian in the language of cycle covers In the case of cubic graphs we have

$$\text{Pfaffian } A = \sum_{c \in cc(G)} \operatorname{sgn}_G(\bar{c}) w(\bar{c}). \tag{6.6}$$

The Pfaffian is always polynomial time computable—no assumptions are necessary except that A is skew symmetric. On the other hand, Theorem 5.1 implies that computing the undirected determinant

$$\text{u-det } G = \sum_{c \in cc(G)} (-1)^{|c|} w(c) \tag{6.7}$$

is #P-complete even for cubic planar graphs. The $(-1)^n$ factor disappears since a cubic graph has an even number of vertices.

When comparing (6.6) and (6.7), we see that the weights $w(c)$ and $w(\bar{c})$ are conveniently related by formula (6.5). In general, the relation between $\operatorname{sgn}_G(\bar{c})$ and $(-1)^{|c|}$ is more complicated; however, for some graphs and their orientations, these are nicely related by Proposition 6.11 below.

Every cycle c in a planar graph G yields a decomposition of the plane P into two closed subsets, the bounded one P_* and the unbounded one P_∞. We have $P = P_* \cup P_\infty$ and $c = P_* \cap P_\infty$. Let $G_* = G \cap P_*$ and $G_\infty = G \cap P_\infty$. Let $v \in c$ be a vertex. We call v an *in-vertex* if the unique edge adjacent to v but not in c belongs to G_*. Otherwise, we call v an *out-vertex*. If c is even then it is bipartite as a subgraph; let V_1, V_2 be the bipartition of its vertices.

Definition 6.8 The *tension* of an even cycle in a cubic planar graph is the absolute value of the difference between the numbers of out-vertices in V_1 and in V_2.

Note that the tension is independent of whether we use the out-vertices or in-vertices in the definition above.

Definition 6.9 An undirected cubic planar graph G is *without tension* if the tension of every even central cycle in G is null.

Remark 6.10 The graphs mentioned in Remark 6.4, those with at most two faces which are bounded by an odd number of edges, are without tension.

Proof Let c be a central cycle in G. Since it has even length, the number of odd faces in P_* must be even. By symmetry between P_*, and P_∞ we may assume that P_∞ has only even faces, including the unbounded one. Therefore P_∞ is bipartite. Since c is central we see that $P_\infty \backslash c$ has a perfect matching and therefore each side of its bipartition has the same number of vertices; denote it by m. Let n_1 and n_2 denote the numbers of out-vertices in each of the two sides of the bipartition of c. We look at P_∞ without the edges in c. The sums of the degrees of vertices in each side of the bipartition must be the same, hence $n_1 + m = n_2 + m$ and therefore the tension $|n_1 - n_2|$ is null. □

Proposition 6.11 *If G is an undirected cubic planar graph without tension, equipped with a semi-Pfaffian orientation, then the function*

$$f_G : cc(G) \longrightarrow \{-1, 1\},$$
$$f_G(c) = (-1)^{|c|} \operatorname{sgn}_G(\bar{c}),$$

is constant.

Proof It is enough to show that for every $c, d \in cc(G)$ we have

$$f_G(c) f_G(d) = 1. \tag{6.12}$$

It is a standard observation that, in a cubic graph, any two cycle covers c and d can be connected by a sequence $c_k \in cc(G)$ with

$$c = c_0, c_1, \ldots, c_r = d$$

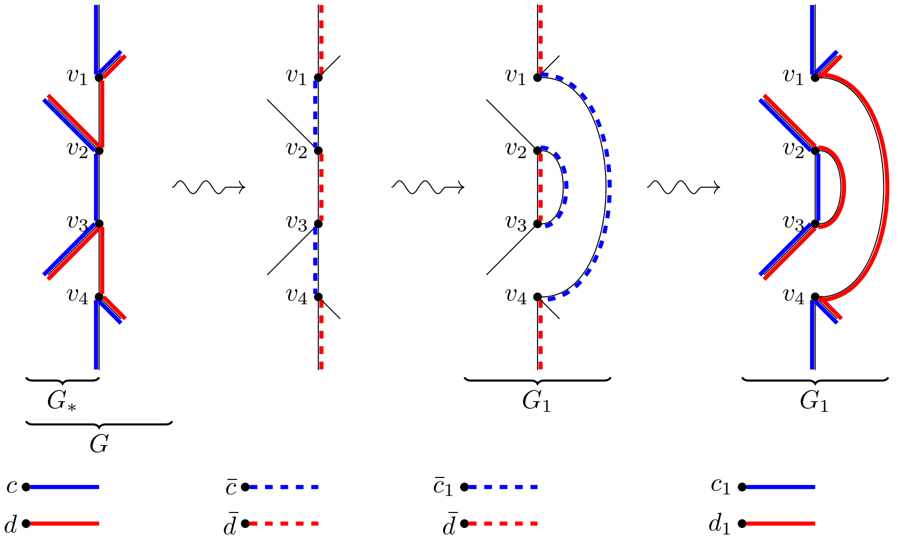

Fig. 18 The modification of G that shortens t and preserves the product $f_G(c)f_G(d) = f_{G_1}(c_1)f_{G_1}(d_1)$

such that for each $k = 0, 1, \ldots, r-1$, the union $\bar{c}_k \cup \bar{c}_{k+1}$ contains exactly one cycle. The cycle cover $\bar{c}_{k+1}$ is obtained from $\bar{c}_k$ by replacing those edges in this cycle which belong to $\bar{c}_k$ with the remaining edges of this cycle—those remaining edges belong to $\bar{d}$.

Thus from now on we may assume that $\bar{c} \cup \bar{d}$ contains only one cycle, or equivalently, that the symmetric difference $c \triangle d$ is a cycle. By assumption, the unique cycle in $\bar{c} \cup \bar{d}$ (or equivalently $c \triangle d$) above has null tension.

The idea of the proof is to construct recursively a series of modifications of the cycle covers c, d, and of the ambient graph G. We construct sequences c_i, d_i and G_i, $i = 0, 1, \ldots, s$, where $G_0 = G$ and $c_i, d_i \in cc(G_i)$, $c_0 = c$, $d_0 = d$. The graphs G_i for $i > 0$ are going to be multigraphs – a pair of vertices may be connected by two edges. At each stage we will have

$$f_{G_i}(c_i) f_{G_i}(d_i) = 1. \tag{6.13}$$

The number of edges in $c_i \triangle d_i$ will be the same as in $c_{i+1} \triangle d_{i+1}$. Both will have no loops and at most one cycle of length exceeding 2. The longest cycle in $c_{i+1} \triangle d_{i+1}$ will be two edges shorter than the longest one in $c_i \triangle d_i$. The $c_{i+1} \triangle d_{i+1}$ will retain the semi-Pfaffian orientation and null tension. Eventually we will have $c_s = d_s$, up to choice of an edge of a multigraph between the same vertices.

A single modification is done in three steps which are shown in Fig. 18. Below we describe such a modification. To simplify the notation we omit the index i and write c_1, d_1, G_1 for $c_{i+1}, d_{i+1}, G_{i+1}$.

We have seen above that we may assume that $t = c \triangle d$ is a single cycle of even length. After modifications $c \triangle d$ will have a unique component of length greater than 2. Since, by assumption, the cycle t has null tension, it must contain two adjacent

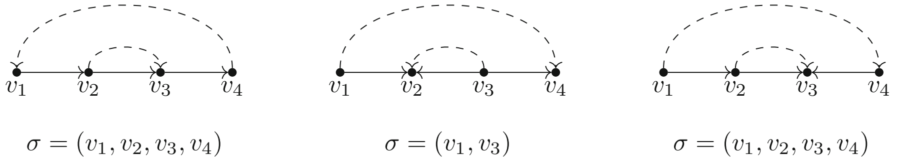

Fig. 19 Possible orientations, up to symmetry, and the corresponding permutations σ

in-vertices or two adjacent out-vertices. We consider the first case, the other being analogous.

The modification is shown in Fig. 18. Below we outline the effect of such modification on the objects we are interested in. The modification:

(1) Changes one of the cycle covers, either c or d. We obtain c_1 and d_1 where either $c_1 = c$ or $d_1 = d$.
(2) Shortens the cycle t by two edges. The new cycle t_1, a connected component of $c_1 \triangle d_1$, retains semi-Pfaffian orientation and null tension.
(3) Transforms the ambient graph G into G_1 so that the two adjacent in-vertices, denoted v_2 and v_3 in Fig. 18, are connected by two parallel edges in G_1.
(4) Leaves the product (6.12) unchanged, that is, $f_G(c) f_G(d) = f_{G_1}(c_1) f_{G_1}(d_1)$.

In the leftmost part of Fig. 18 we see the adjacent in-vertices v_2 and v_3 with the surrounding fragments of the cycle covers c and d. Possibly swapping the names c and d, we may assume that the edge $\{v_2, v_3\}$ belongs to c, as shown in Fig. 18. The vertices v_1 and v_4 may be, independently, either in-vertices or out-vertices.

In the second part of Fig. 18 we mark the complements $\bar{c}$ and $\bar{d}$.

In the third part we apply a permutation σ to the vertices $\{v_1, v_2, v_3, v_4\}$. The permutation moves the edges $\{v_1, v_2\}$ and $\{v_3, v_4\}$, together with their orientations, and is subject to the following conditions:

(1) $\sigma(\{v_1, v_2\}) = \{v_2, v_3\}$ so that the old edge $\{v_2, v_3\}$ and the new edge $\sigma(\{v_1, v_2\})$ have the same orientation.
(2) $\sigma(\{v_3, v_4\}) = \{v_1, v_4\}$ so that the induced orientation of $\{v_1, v_4\}$ agrees with the direction of traversal around t which is induced by the orientations of the even number of the original edges $\{v_1, v_2\}$, $\{v_2, v_3\}$ and $\{v_3, v_4\}$.

This modification removes two vertices v_2 and v_3 from the cycle t which reduces its length by 2, from some $2k$ to $2(k-1)$. Condition (2) implies that the parity of the number of edges directed clockwise (as well as those directed counterclockwise) changes, therefore the modified cycle t_1 retains the semi-Pfaffian orientation. Figure 19 lists, up to symmetry, all the possible original orientations of $\{v_1, v_2\}$, $\{v_2, v_3\}$ and $\{v_3, v_4\}$, and the corresponding permutations σ.

Since in every case shown in Fig. 19, the permutation σ is odd, it changes the sign (see (6.1)) of the perfect matching $\bar{c}$, viewed as a summand of the Pfaffian, in (6.2). This way we obtain

$$\operatorname{sgn}_{G_1}(\bar{c}_1) = -\operatorname{sgn}_G(\bar{c}). \tag{6.14}$$

In the rightmost part of Fig. 18 we see the cycle covers c_1 and d_1, the complements of $\bar{c}_1$ and $\bar{d}$ in the new graph G_1. We see that $c_1 = c$, but d_1 is different from d.

The modification shown in Fig. 18 changes the parity of the number of cycles in the cycle cover d so that

$$(-1)^{|d_1|} = -(-1)^{|d|}. \tag{6.15}$$

Combining (6.14) and (6.15) we obtain

$$\begin{aligned} f_{G_1}(c_1) f_{G_1}(d_1) &= (-1)^{|c_1|} \operatorname{sgn}_{G_1}(\bar{c}_1)(-1)^{|d_1|} \operatorname{sgn}_{G_1}(\bar{d}_1) \\ &= (-1)^{|c|}(-\operatorname{sgn}_G(\bar{c}))(-(-1)^{|d|}) \operatorname{sgn}_G(\bar{d}) \\ &= f_G(c) f_G(d), \end{aligned} \tag{6.16}$$

thus the modification does not change the product.

We have deleted two in-vertices from t, one from either part of the bipartition of t, so that the modified cycle t_1 retains the null tension.

We now repeat the operation shown in Fig. 18; at every step the length of t is reduced by 2, until at some step s its length is 2. At that point we see that every edge in $\bar{c}_s$ corresponds to an edge in $\bar{d}_s$, which has the same end points. This correspondence preserves the orientations. Therefore $\operatorname{sgn}_{G_s}(\bar{c}_s) = \operatorname{sgn}_{G_s}(\bar{d}_s)$ and similarly $|c_s| = |d_s|$, so that we obtain

$$\begin{aligned} f_G(c) f_G(d) &= \cdots = f_{G_s}(c_s) f_{G_s}(d_s) \\ &= (-1)^{|c_s|} \operatorname{sgn}_{G_s}(\bar{c}_s)(-1)^{|d_s|} \operatorname{sgn}_{G_s}(\bar{d}_s) \\ &= 1 \end{aligned} \tag{6.17}$$

where the dots indicate a sequence of equations given by (6.16). This completes the proof. □

As a corollary, we obtain the following.

Theorem 6.18 *If G is a weighted undirected cubic planar graph with a semi-Pfaffian orientation and without tension, and all the weights in G are nonzero, then, up to sign,*

$$\text{u-det}\, G = p \cdot \text{Pfaffian}(A_{inv}),$$

where p is the product of the weights of all the edges in G and A_{inv} is the matrix whose nonzero entries are the inverses of the nonzero entries of the skew symmetric adjacency matrix of G with the orientation.

Proof This is an immediate consequence of Proposition 6.11 and the identities (6.5), (6.6) and (6.7). □

Funding The authors have not disclosed any funding.

Data Availability Statement Data sharing is not applicable to this article as no datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Appendix

The constructions presented in the preceding sections were chosen so as to minimize the number of edges and labels. In this section we prove

Theorem 7.1 *Computing the undirected determinant of cubic planar graphs is #P-complete even when we restrict to graphs whose edges all have weight* 1.

Proof The first place in the proof of Theorem 5.1 where labels different from 1 and -1 appear is Fig. 16. There we have the degree 4 vertex gadget whose four edges have weight $-\frac{1}{2}$. If we want to have only integral weights, we multiply the labels of all the inner edges of this gadget by 2. The signature of the original gadget is shown in Fig. 17. Doubling the weights of the edges does not affect the value 0 of the signature, while the value -4 is multiplied by 2^{47}—in that case every cycle cover meets the gadget in 47 internal edges and 2 external ones. Therefore, if the graph has v degree 4 vertex gadgets and each is modified as above, the undirected determinant is multiplied by $(2^{47})^v$, and we conclude that computing the undirected determinant of cubic planar graphs, with edge weights in the set $\{-2, -1, 1, 2\}$ is $\#P$-complete.

It remains to construct determinantal gadgets whose edges all have weight 1 and whose signature is that of a single edge of integral weight. The gadget of weight 1 is just a single edge. The gadget replacing an edge of weight -1 is shown in Fig. 20.

Given two gadgets replacing edges of weights a and b, the gadget for an edge of weight ab is shown in Fig. 21. Computing of its signature is straightforward.

We also have an $a + b$ edge gadget shown in Fig. 22. Note that the null edge gadget defined in Fig. 15 includes the -1 edge gadget defined in Fig. 20.

In the first paragraph of this proof we have proved that computing the undirected determinant of cubic planar graphs, with edge weights in the set $\{-2, -1, 1, 2\}$, is $\#P$-complete. Now we can replace these edges with suitable gadgets having weight 1 edges only. This completes the proof. □

Using binary expansion we can replace an edge labeled with any integer by a suitable composition of gadgets constructed in the proof above. Thus we obtain the following.

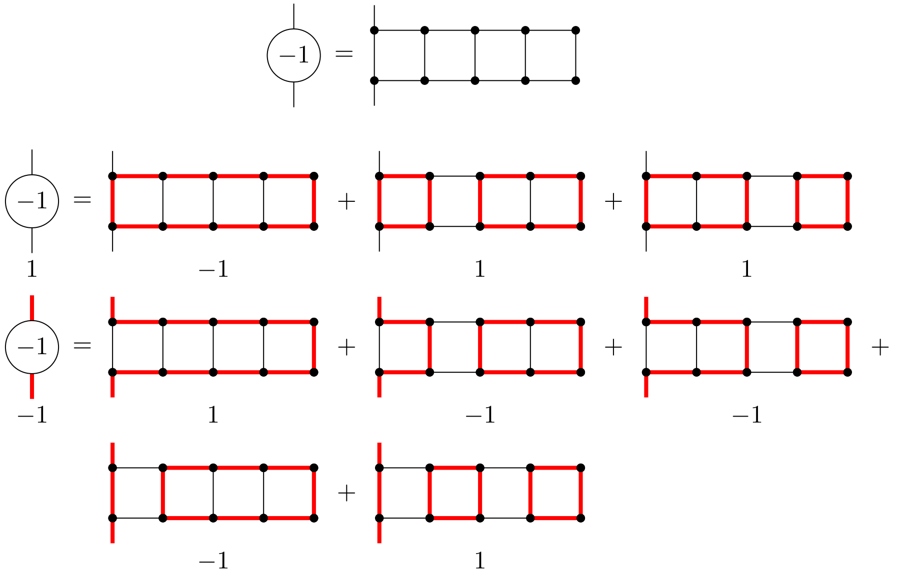

Fig. 20 The -1 edge gadget and its signature

Fig. 21 The ab edge gadget

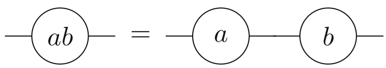

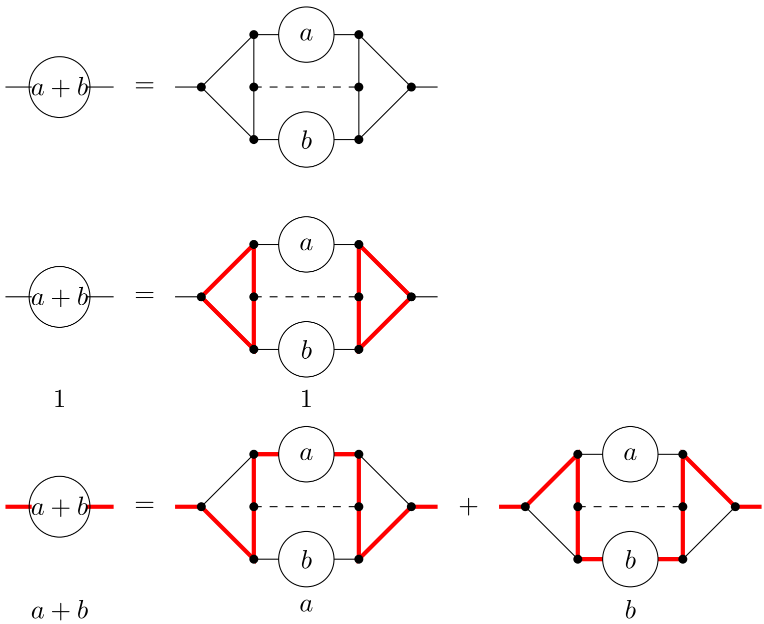

Fig. 22 The $a+b$ edge gadget and its signature

Corollary 7.2 *Every cubic planar graph with integral edge weights can be replaced with a cubic planar graph with edge weights* 1*. The increase of the size of the graph is logarithmic in the size of the original edge weights.*

Obviously, in the case of the undirected permanent it is impossible to construct a weight 1 gadget that would replace an edge of negative weight.

References

1. Cai, J.-Y., Fu, Z., Guo, H., Williams, T.: A Holant dichotomy: is the FKT algorithm universal? 2015 IEEE 56th Annual Symposium on Foundations of Computer Science-FOCS 2015, pp. 1259–1276. IEEE Computer Soc., Los Alamitos (2015)
2. Cai, J.-Y., Gorenstein, A.: Matchgates revisited. Theory Comput. **10**, 167–197 (2014)
3. Chien, S., Harsha, P., Sinclair, A., Srinivasan, S.: Almost settling the hardness of noncommutative determinant. In: STOC'11-Proceedings of the 43rd ACM Symposium on Theory of Computing, pp. 499–508. ACM, New York (2011)
4. de Rugy-Altherre, N.: Determinant versus permanent: salvation via generalization? In: The Nature of Computation, Lecture Notes in Comput. Sci., vol. 7921, pp. 87–96. Springer, Heidelberg (2013)
5. Kasteleyn, P.W.: The statistics of dimers on a lattice: I. The number of dimer arrangements on a quadratic lattice. Physica **27**, 1209–1225 (1961)
6. Kasteleyn, P.W.: Graph theory and crystal physics. In: Graph Theory and Theoretical Physics, pp. 43–110. Academic Press, London (1967)
7. Mahajan, M., Vinay, V.: Determinant: combinatorics, algorithms, and complexity. Chicago J. Theoret. Comput. Sci. Article 5, 26 (1997)
8. Mertens, S., Moore, C.: The complexity of the fermionant and immanants of constant width. Theory Comput. **9**, 273–282 (2013)
9. Rote, G.: Division-free algorithms for the determinant and the Pfaffian: algebraic and combinatorial approaches. In: Computational discrete mathematics, Lecture Notes in Comput. Sci., vol. 2122, pp. 119–135. Springer, Berlin (2001)
10. Temperley, H.N., Fisher, M.E.: Dimer problem in statistical mechanics-an exact result. Philos. Mag. **6**, 1061–1063 (1961)
11. Thomas, R.: A survey of Pfaffian orientations of graphs. International Congress of Mathematicians. Vol. III, pp. 963–984. Eur. Math. Soc., Zürich (2006)
12. Valiant, L.G.: Expressiveness of matchgates. Theoret. Comput. Sci. **289**, 457–471 (2002)
13. Valiant, L.G.: Quantum circuits that can be simulated classically in polynomial time. SIAM J. Comput. **31**, 1229–1254 (2002)
14. Valiant, L.G.: Accidental algorithms. In: 2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06), Berkeley, pp. 509–517 (2006)
15. Valiant, L.G.: Holographic algorithms. SIAM J. Comput. **37**, 1565–1594 (2008)