2015 IEEE 56th Annual Symposium on Foundations of Computer Science

Parameterizing the Permanent: genus, apices, minors, evaluation mod 2^k

Radu Curticapean
Institute for Computer Science and Control,
Hungarian Academy of Sciences (MTA SZTAKI)
Budapest, Hungary
curticapean@cs.uni-sb.de

Mingji Xia
State Key Laboratory of Computer Science,
Institute of Software, Chinese Academy of Sciences.
Beijing, China
mingji@ios.ac.cn

Abstract

We identify and study relevant structural parameters for the problem $\mathrm{PerfMatch}$ of counting perfect matchings in a given input graph G. These generalize the well-known tractable planar case, and they include the *genus* of G, its *apex number* (the minimum number of vertices whose removal renders G planar), and its *Hadwiger number* (the size of a largest clique minor).

To study these parameters, we first introduce the notion of *combined matchgates*, a general technique that bridges parameterized counting problems and the theory of so-called Holants and matchgates: Using combined matchgates, we can simulate certain non-existing gadgets F as linear combinations of $t = \mathcal{O}(1)$ existing gadgets. If a graph G features k occurrences of F, we can then reduce G to t^k graphs that feature only existing gadgets, thus enabling parameterized reductions.

As applications of this technique, we simplify known $4^g n^{O(1)}$ time algorithms for $\mathrm{PerfMatch}$ on graphs of genus g. Orthogonally to this, we show $\#\mathsf{W}[1]$-hardness of the permanent on k-apex graphs, implying its $\#\mathsf{W}[1]$-hardness under the Hadwiger number. Additionally, we rule out $n^{o(k/\log k)}$ time algorithms under the counting exponential-time hypothesis $\#\mathsf{ETH}$.

Finally, we use combined matchgates to prove $\oplus\mathsf{W}[1]$-hardness of evaluating the permanent modulo 2^k, complementing an $\mathcal{O}(n^{4k-3})$ time algorithm by Valiant and answering an open question of Björklund. We also obtain a lower bound of $n^{\Omega(k/\log k)}$ under the parity version $\oplus\mathsf{ETH}$ of the exponential-time hypothesis.

Index Terms

permanent; perfect matchings; parameterized counting complexity; genus; apex number; graph minors; Hadwiger number; modular counting complexity; matchgates

I. Introduction

The study of counting problems has become a classical subfield of computational complexity since Valiant's seminal papers [1], [2] that introduced the class $\#\mathsf{P}$ and established $\#\mathsf{P}$-completeness of computing (for graphs G with edge-weights $w : E(G) \to \mathbb{Q}$) the quantity

$$\mathrm{PerfMatch}(G) := \sum_{\substack{M \subseteq E(G) \\ \text{perfect matching of } G}} \prod_{e \in M} w(e).$$

In statistical physics, $\mathrm{PerfMatch}$ is known as the *partition function* of the *dimer model* [3], [4], [5], and the first nontrivial algorithms for its evaluation stem from this area. This includes the celebrated *FKT method*, a polynomial-time algorithm for computing $\mathrm{PerfMatch}$ on planar graphs [5]. Roughly speaking, this algorithm proceeds as follows: Given a planar graph G, it constructs a *Pfaffian orientation* F of G, which we may view as a subset $F \subseteq E(G)$ with the following miraculous property: If we define a matrix A from the adjacency matrix of G by flipping the signs of edges in F, then $(\mathrm{PerfMatch}(G))^2 = \det(A)$. Overall, this yields a reduction from planar $\mathrm{PerfMatch}$ to the determinant.

In algebra and combinatorics, the quantity $\mathrm{PerfMatch}(G)$ for a bipartite graph G with $n + n$ vertices is better known as the *permanent* of the biadjacency matrix A of G, defined by

$$\mathrm{perm}(A) = \sum_{\substack{\sigma:[n]\to[n] \\ \text{is a permutation}}} \prod_{i=1}^{n} A_{i,\sigma(i)}.$$

This quantity is central to algebraic complexity theory, which aims at proving the permanent to be inherently harder than the similar-looking determinant [6], [7], [8]. This would imply an algebraic analogue of $\mathsf{P} \neq \mathsf{NP}$ [9].

In order to obtain a more refined view on the complexity of the permanent, and to cope with its hardness in view of practical applications, various relaxations of this problem were studied: A celebrated randomized approximation scheme [10], [11] allows one to *approximate* the permanent on matrices with non-negative entries. Furthermore, on some *restricted* graph classes, $\mathrm{PerfMatch}$ can be solved in time $\mathcal{O}(n^3)$: This includes the above-mentioned planar graphs, and in fact, all graph classes of bounded genus [12], [13], [14]. As another relaxation, it was shown in Valiant's original paper [1] that the permanent *modulo* $m = 2^k$ can be computed in time $n^{\mathcal{O}(k)}$, but for all $m \neq 2^k$, it is NP-hard under randomized reductions.

0272-5428/15 $31.00 © 2015 IEEE
DOI 10.1109/FOCS.2015.65

In this paper, we consider another such refinement (and combine existing ones) by investigating the permanent in the framework of *parameterized complexity*. This area was initiated by Downey and Fellows [15], [16] and was adapted to counting problems by Flum and Grohe [17] and McCartin [18]. In parameterized counting complexity, the objects in study are counting problems that come with *parameterizations* $\pi : \{0,1\}^* \to \mathbb{N}$, and a central question is whether such problems are *fixed-parameter tractable (fpt)*: A given problem is fpt if it can be solved in time $f(\pi(x))|x|^{\mathcal{O}(1)}$ on input x, for a computable function f that depends only on the parameter value, but not on $|x|$. We can also give evidence that problems are *not* fpt by proving their $\#\mathsf{W}[1]$-hardness, the parameterized analogue of $\#\mathsf{P}$-hardness. For more details, consider Section II.

By studying natural parameterizations π of the input, we obtain a fine-grained complexity analysis that could not be achieved by considering the input size $|x|$ alone. For instance, consider the problem $\mathrm{VertexCover}$, which asks whether a graph G on n vertices admits a vertex-cover of size k. This problem is NP-complete, but it can be solved in time $n^{\mathcal{O}(k)}$ for every fixed k, and it is actually even fpt in the parameter k, as we can find [15] and even count [19] vertex-covers of size k in time $2^k n^{\mathcal{O}(1)}$. On the other hand, we can decide in polynomial time whether G contains a matching of size k, but the problem of counting k-matchings is $\#\mathsf{P}$-complete, and in fact even $\#\mathsf{W}[1]$-complete when parameterized by k [20], [21].

A. Genus, apices and excluded minors

To investigate the parameterized complexity of the permanent, we first identify interesting parameterizations for this problem. For instance, the maximum degree $\Delta(G)$ of the input graph G, is not particularly interesting, since the permanent is already $\#\mathsf{P}$-complete on 3-regular graphs [22]. That is, even an $n^{f(\Delta(G))}$ time algorithm (and an fpt-algorithm in particular) would imply $\mathsf{P} = \#\mathsf{P}$. However, it turns out that the known polynomial-time solvable graph classes for $\mathrm{PerfMatch}$ point us towards a natural parameter, namely the size of a smallest excluded minor. (A minor H of a graph G can be obtained by deletions of edges and/or vertices, and contraction of edges.) To explain this, we survey the known algorithms for $\mathrm{PerfMatch}$.

- Excluding $K_{3,3}$ or K_5: It was shown by Little [23] and Vazirani [24] that $\mathrm{PerfMatch}$ can be solved in time $\mathcal{O}(n^3)$ on graphs excluding the minor $K_{3,3}$. A similar result was recently shown by Straub et al. [25] for graphs excluding K_5. Note that the FKT method gives an $\mathcal{O}(n^3)$ time algorithm on graphs excluding *both* $K_{3,3}$ and K_5, whereas the two above algorithms show that excluding *either* minor entails the polynomial-time solvability of $\mathrm{PerfMatch}$. For the $K_{3,3}$-free case, this is shown by constructing a Pfaffian orientation. The K_5-free case was shown by a different technique; in particular, K_5-free graphs do not necessarily admit Pfaffian orientations.
- Excluding single-crossing minors: Extending the above item, it was recently shown by Curticapean [26] that $\mathrm{PerfMatch}$ can be solved in time $\mathcal{O}(n^4)$ on any class excluding a fixed *single-crossing minor* H, i.e., a minor that can be drawn in the plane with at most one crossing, such as $K_{3,3}$ or K_5. In fact, it is shown that $\mathrm{PerfMatch}$ is fpt in the size of the smallest excluded single-crossing minor. This algorithm does not inherently rely upon Pfaffian orientations, apart from a black-box algorithm for planar $\mathrm{PerfMatch}$.
- Bounded-genus graphs: Another line of extensions of the FKT method is to graphs of bounded *genus*: It was shown independently by Gallucio and Loebl [12], Tesler [13] and Regge and Zechina [14] that $\mathrm{PerfMatch}$ can be solved in time $\mathcal{O}(4^g n^3)$ on n-vertex graphs G of genus g, so $\mathrm{PerfMatch}$ is fpt when parameterized by the genus of G. All algorithms proceed by expressing $\mathrm{PerfMatch}(G)$ as the linear combination of 4^g determinants derived from Pfaffian orientations. In the present paper, we give an alternative proof of this theorem that proceeds by reduction to 4^g instances of planar $\mathrm{PerfMatch}$. Together with the previous item, this eliminates the need for Pfaffian orientations from all known algorithms for $\mathrm{PerfMatch}$ except for the planar case.

We are ready to draw the following conclusion: Every *known* polynomial-time solvable graph class for $\mathrm{PerfMatch}$ excludes some fixed minor.[1] This is clear for the first two items, and furthermore, the graphs of genus $g \in \mathbb{N}$ are easily seen to exclude a complete graph $K_{\mathcal{O}(g)}$. Since this shows that excluded minors have been a driving force behind polynomial-time algorithms for $\mathrm{PerfMatch}$, it is natural to study this problem under the more general *Hadwiger number*

$$\mathrm{hadw}(G) := \max\{k \in \mathbb{N} : \ G \text{ contains a } K_k \text{ minor}\}.$$

Note that planar graphs have Hadwiger number at most 4. More generally, if the genus of G or the size of the smallest excluded single-crossing minor is bounded, then $\mathrm{hadw}(G)$ is bounded as well, but the converse does not hold. However, the *Graph Structure Theorem* [27], a celebrated result in graph minor theory [28], yields a decomposition of the graphs with fixed Hadwiger number k into graphs that have genus $c = c(k)$ except for c occurrences of certain defects, namely so-called vortices and apices. Such decompositions have proven immensely useful for fpt-algorithms on graphs excluding fixed minors, see [29], [30], [31], [32], [33], [34]. If a problem can be solved efficiently on planar instances and we can extend this to bounded-genus instances, as in the case of $\mathrm{PerfMatch}$, then with a leap of faith, the Graph Structure Theorem allows us to hope for an

[1] This statement comes with a caveat: For instance, we can determine the number of perfect matchings in a complete graph in polynomial time (by a closed formula). The class of complete graphs clearly excludes no fixed minor. However, we cannot solve the (weighted) problem $\mathrm{PerfMatch}$ on this class in polynomial time, as edge-weights would allow us to simulate arbitrary graphs, for which counting perfect matchings is $\#\mathsf{P}$-complete.

fpt-algorithm under the more general Hadwiger number. Our following negative result however shatters these hopes for the case of $\mathrm{PerfMatch}$.

Theorem I.1. *The zero-one permanent is* $\#\mathsf{W}[1]$*-hard when parameterized by the Hadwiger number. In other words, computing* $\mathrm{PerfMatch}$ *is* $\#\mathsf{W}[1]$*-hard when parameterized by the Hadwiger number, even on bipartite graphs without edge-weights.*

We show this by proving the following stronger statement: Let us define the apex number

$$\mathrm{apex}(G) := \min\{k \in \mathbb{N} \mid \exists S \subseteq V(G) \text{ of size } k : \ G - S \text{ is planar}\}.$$

This parameter, studied in [35], measures the distance of a graph to planarity by vertex deletions. Note that planar graphs have apex number 0. Using the apex number as parameter, we can generalize planar graphs in a way that is orthogonal to the genus parameter: There are graphs on which any one of these parameters is bounded, while the other is not. However, it can be verified that $\mathrm{hadw}(G) \le \mathcal{O}(\mathrm{apex}(G))$. This allows us to obtain Theorem I.1 as a corollary from the following result, which we consider to be of independent interest.

Theorem I.2. *The permanent is* $\#\mathsf{W}[1]$*-hard when parameterized by the apex number. Assuming the exponential-time hypothesis* $\#\mathsf{ETH}$*, it admits no* $n^{o(k/\log k)}$ *time algorithm on* k*-apex graphs with* n *vertices.*

This contrasts with the fpt-algorithm for $\mathrm{PerfMatch}$ when parameterized by genus. We observe that $\mathrm{PerfMatch}$ can be computed easily in time $n^{k+\mathcal{O}(1)}$ on k-apex graphs by means of brute-force, so the lower bound under $\#\mathsf{ETH}$ is almost tight. However, it should be noted that no similar algorithm is known for the Hadwiger number: At least to us, it remains an important open question whether $\mathrm{PerfMatch}$ can be solved in time $n^{f(k)}$ on graphs excluding the complete graph K_k.

B. Evaluating the permanent modulo 2^k

In the following, we depart from structural parameters of the input graph G and consider the evaluation of the permanent modulo 2^k. In the seminal paper [1], not only did Valiant prove $\#\mathsf{P}$-completeness of the permanent, but he also studied the complexity of evaluating the permanent modulo fixed numbers $m \in \mathbb{N}$.

Observe that $\mathrm{perm}(A)$ and $\det(A)$ are equivalent modulo 2 for any matrix A, giving a polynomial-time algorithm for the permanent modulo 2. On the other hand, for odd primes p, Valiant's original proof shows that the permanent modulo p is $\mathsf{Mod}_p\mathsf{P}$-complete. That is, we can reduce counting satisfying assignments to 3-CNF formulas modulo p to the permanent modulo p. This also shows its NP-hardness under randomized reductions, and this holds more generally whenever the modulus m is not a power of two.

For the remaining cases $m = 2^k$ however, Valiant [1] showed an $\mathcal{O}(n^{4k})$ time algorithm for evaluating the permanent modulo 2^k on n-vertex graphs, which was recently improved to $\mathcal{O}(n^k)$ by Björklund, Husfeldt and Lyckberg [36]. Given these results, it is natural to study this problem in the framework of parameterized complexity, thus asking whether we can compute the permanent modulo 2^k in time $n^{o(k)}$ or even $f(k)n^{\mathcal{O}(1)}$. This was also posed as an open problem in [36]. Please recall that this question is indeed only interesting for $m = 2^k$: As stated in the previous paragraph, on all other *fixed* $m \in \mathbb{N}$, the problem is NP-hard under randomized reductions.

We rule out the fixed-parameter tractability by the following stronger hardness result, which also establishes an unexpected connection to the apex parameter introduced before: Evaluating the permanent modulo 2^k on k-apex graphs is $\oplus\mathsf{W}[1]$-hard, that is, an fpt-algorithm for this problem would imply one for counting k-cliques modulo 2. This problem however is $\mathsf{W}[1]$-hard under randomized reductions by a recent result of Björkund, Dell and Husfeldt [37]. We also obtain an almost-tight lower bound under $\oplus\mathsf{ETH}$, the parity version of the exponential-time hypothesis ETH.

Theorem I.3. *The evaluation of the permanent modulo* 2^k *is* $\oplus\mathsf{W}[1]$*-hard when parameterized by* k*, even when restricted to* k*-apex graphs. Assuming* $\oplus\mathsf{ETH}$*, there is no* $n^{o(k/\log k)}$ *time algorithm for this problem.*

We prove Theorem I.3 by reduction from the following problem $\oplus\mathrm{PartitionedSub}$: Given vertex-colored graphs H and G as input, where each color in H appears exactly once, count modulo 2 the subgraphs of G that are isomorphic to H, respecting colors. It was shown that the decision version of this problem can be reduced to $\oplus\mathrm{PartitionedSub}$ by means of randomized reductions [37]. Furthermore, assuming $\oplus\mathsf{ETH}$, an argument by Marx [38] implies that $\oplus\mathrm{PartitionedSub}$ cannot be solved in time $n^{o(\ell/\log \ell)}$ for ℓ-edge graphs H and n-vertex graphs G.

In our reduction, we transform a given instance (H, G) for $\oplus\mathrm{PartitionedSub}$ with an ℓ-edge graph H to 3^ℓ instances of the permanent modulo $2^{2\ell+1}$ on 2ℓ-apex graphs with $\mathcal{O}(\ell^2 n^2)$ vertices. Thus, if we can prove better lower bounds for finding k-edge subgraphs, then those bounds carry over to the seemingly unrelated problem of evaluating permanents modulo 2^k, even on k-apex graphs. On the other hand, a randomized $n^{o(k)}$ time algorithm for the permanent modulo 2^k on k-apex graphs would imply one for $\mathrm{PartitionedSub}$ on k-edge graphs H, thus falsifying a hypothesis posed by Marx [38].

C. Proof technique: Linear combinations of signatures

We phrase our proofs in the language of so-called Holant problems [39] and matchgates [39], [40], [41]. Due to space limitations, we refer to Section III for an introduction into this topic. In our proofs, we reformulate parameterized counting problems as Holant problems (specific weighted sums over assignments to the edges of graphs) and then try to realize the occurring signatures (local constraints at vertices) by certain matchgates (gadgets, graph fragments). However, many required signatures cannot be realized by matchgates. The key technical idea underlying our paper is that such unrealizable signatures can sometimes still be realized as *linear combinations* of matchgate signatures.

To this end, we proceed as follows: First, we show how to simulate non-existing gadgets F as the linear combination of realizable gadgets $F_1, \ldots, F_t$, typically with $t = \mathcal{O}(1)$. Then, if a graph G features k occurrences of F, we can easily reduce G to t^k graphs that feature only occurrences of $F_1, \ldots, F_t$. Each of these t^k graphs can then be handled by an algorithm (when we aim at positive results) or by an oracle call (when proving hardness results). The generality of our technique allows it to be applied to various parameterized problems. For instance, a recent $\#\mathsf{W}[1]$-hardness proof for counting k-matchings [21] can also be rephrased in this framework.[2]

II. General preliminaries

For $n \in \mathbb{N}$, we write $[n] := \{1, \ldots, n\}$. The graphs G in this paper are undirected, but they may feature parallel edges and edge-weights. All *hardness results* are however shown for *simple* graphs featuring no parallel edges and no edge-weights. We write $uv \in E(G)$ for an edge of G, and given $v \in V(G)$, we denote the edges incident with v by $I(v)$. Sometimes, we consider graphs to be embedded on surfaces, see [44].

For numbers $n \in \mathbb{N}$, we abbreviate $\oplus n := (n \bmod 2)$. Given a bitstring $x \in \{0,1\}^*$, we write $\mathrm{hw}(x) := \sum_i x_i$ for its *Hamming weight*, and we define $\mathtt{ODD}(x) := \oplus\mathrm{hw}(x)$ and $\mathtt{EVEN}(x) := 1 - \oplus\mathrm{hw}(x)$. We write $\mathrm{supp}(f)$ for the support of a function f. For predicates φ, we define

$$[\varphi] := \begin{cases} 1 & \varphi \text{ is true,} \\ 0 & \text{otherwise.} \end{cases}$$

Let A and B be sets; we consider subsets of $A \times B$. For $b \in B$, we write $(\star, b) = \{(a,b) \mid a \in A\}$ for the *column* at b. For $a \in A$, we write $(a, \star) = \{(a,b) \mid b \in B\}$ for the *row* at a. We use this notation only when A and B are clear from the context. For $k \in \mathbb{N}$, we say that $(i,j) \in [k]^2$ and $(i',j') \in [k]^2$ are *vertically adjacent* if $|i - i'| = 1$ and $j = j'$. Likewise, we call such pairs *horizontally adjacent* if $|j - j'| = 1$ and $i = i'$.

A. Parameterized complexity

Parameterized counting problems are problems A/π, where $\mathrm{A} : \{0,1\}^* \to \mathbb{C}$ is a counting problem and $\pi : \{0,1\}^* \to \mathbb{N}$ is a polynomial-time computable parameterization, see [17]. We have $\mathrm{A}/\pi \in \mathsf{FPT}$ if A can be solved in time $f(\pi(x))|x|^{\mathcal{O}(1)}$, and $\mathrm{A}/\pi \in \mathsf{XP}$ if it can be solved in time $|x|^{f(\pi(x))}$, where $f : \mathbb{N} \to \mathbb{N}$ is a computable function. In the following, we define the classes $\mathsf{W}[1]$, $\#\mathsf{W}[1]$ and $\oplus\mathsf{W}[1]$ from the introduction, using the following reduction notions.

Definition II.1 ([17])**.** Let A/π and B/π' be parameterized counting problems.

- We call $f : \{0,1\}^* \to \{0,1\}^*$ a *parsimonious fpt-reduction* and write $\mathrm{A}/\pi \leq_{fpt}^{pars} \mathrm{B}/\pi'$ if there are computable functions r, s such that the following holds for all $x \in \{0,1\}^*$:
  1) We have $\mathrm{A}(x) = \mathrm{B}(f(x))$.
  2) The running time of f is bounded by $r(\pi(x)) \cdot |x|^{\mathcal{O}(1)}$.
  3) We have $\pi'(f(x)) \leq s(\pi(x))$.

  If A and B are decision problems, replace the first condition by "$x \in \mathrm{A}$ iff $f(x) \in \mathrm{B}$", and write $\mathrm{A}/\pi \leq_{fpt} \mathrm{B}/\pi'$.
- We call an algorithm $\mathbb{T}$ a *Turing fpt-reduction* and write $\mathrm{A}/\pi \leq_{fpt}^{T} \mathrm{B}/\pi'$ if there are computable functions r and s such that the following holds for all $x \in \{0,1\}^*$: Firstly, the running time of $\mathbb{T}$ on x is bounded by $r(\pi(x))|x|^{\mathcal{O}(1)}$. Secondly, every oracle query y issued by $\mathbb{T}$ on x satisfies $\pi'(y) \leq s(\pi(x))$.

We introduce $\mathsf{W}[1]$, $\oplus\mathsf{W}[1]$ and $\#\mathsf{W}[1]$ as the closures of canonical clique-related problems under fpt-reductions.

Definition II.2. Consider the following parameterized problems and complexity classes:

- Let Clique/k denote the problem of *deciding*, on input a graph G and $k \in \mathbb{N}$, whether G contains a k-clique. Let $\mathsf{W}[1]$ denote the set of all problems A/π with $\mathrm{A}/\pi \leq_{fpt} \mathrm{Clique}/k$.
- Let $\#\mathrm{Clique}/k$ denote the problem of determining, on input G and k, the *number* of k-cliques in G. Let $\#\mathsf{W}[1]$ denote the set of all problems A/π with $\mathrm{A}/\pi \leq_{fpt}^{pars} \#\mathrm{Clique}/k$.

[2]As pointed out by Tyson Williams, the notion of combined signatures bears some similarities to that of *vanishing signatures* [42], [43]. These however apply linear combinations in a quite different setting. In particular, they consider no connections to parameterized complexity.

- Let $\oplus\mathrm{Clique}/k$ denote the problem of *deciding*, on input G and k, whether G contains an *odd* number of k-cliques. Let $\oplus\mathsf{W}[1]$ denote the set of all A/π with $\mathrm{A}/\pi \leq_{fpt} \oplus\mathrm{Clique}/k$.

It is a standard assumption of parameterized complexity theory that $\mathsf{FPT} \neq \mathsf{W}[1]$ holds, implying $\mathsf{FPT} \neq \#\mathsf{W}[1]$. The problem Clique/k is $\mathsf{W}[1]$-complete by definition, so this assumption can equivalently be considered as the statement that Clique/k is not fixed-parameter tractable. Furthermore, it has been recently shown in [37, Theorem 5] that $\oplus\mathrm{Clique}/k$ is $\mathsf{W}[1]$-hard under randomized parameterized reductions with constant one-sided error. Therefore, an fpt-algorithm for $\oplus\mathrm{Clique}/k$ would imply a randomized fpt-algorithm for Clique/k, which is considered almost as unlikely as $\mathsf{FPT} = \mathsf{W}[1]$.

B. Exponential-time complexity

We also consider conditional lower bounds on the running times required to solve problems. These are based on different exponential-time hypotheses, introduced by [45], [46] and [47].

Definition II.3. The exponential-time hypothesis ETH, introduced in [45], [46], claims that the satisfiability of 3-CNF formulas on n variables cannot be decided in time $2^{o(n)}n^{O(1)}$. The hypothesis $\#\mathsf{ETH}$ postulates the same lower bound for counting 3-CNF formulas, and $\oplus\mathsf{ETH}$ postulates the same for computing the parity of satisfying assignments.

The hypothesis ETH implies a lower bound for Clique/k, and thus also $\mathsf{FPT} \neq \mathsf{W}[1]$: It was shown in [48], [49] that Clique/k cannot be solved in time $n^{o(k)}$ on n-vertex graphs. Furthermore, if a problem A/π admits a lower bound of $n^{g(k)}$ under ETH, and we can reduce A/π to B/π' with a reduction f that satisfies $\pi'(f(x)) \in \mathcal{O}(\pi(x))$ for all x, then it can be seen that B/π' admits a lower bound of $n^{\Omega(g(k))}$ under ETH.

By an isolation argument similar to the Valiant-Vazirani theorem [50], it was shown in [51] that a $2^{o(n)}$ time algorithm for counting satisfying assignments to 3-CNF formulas modulo 2 implies a randomized $2^{o(n)}$ time algorithm for deciding the existence of a satisfying assignment. In other words, a randomized version rETH of ETH implies $\oplus\mathsf{ETH}$, see also [47].

C. The complexity of grid tilings

We will reduce from the problem of counting grid tilings, possibly modulo two. The decision version of this problem was introduced by Marx [52] in order to obtain lower bounds for planar multiway cut, but grid tilings have since proven to be a generally useful reduction source for proving hardness of planar-ish problems [29].

Definition II.4. The inputs to the problem $\mathrm{GridTiling}$ are numbers $n, k \in \mathbb{N}$, together with a set $\mathcal{C} \subseteq [k]^2$ and a function $\mathcal{T}$ that maps from $\mathcal{C}$ into the power-set of $[n]^2$. The task is to decide whether there exists a *grid tiling* of $\mathcal{T}$, i.e., a function $a : [k]^2 \to [n]^2$ such that:

1) For horizontally adjacent $\kappa, \kappa' \in [k]^2$, the first components of $a(\kappa)$ and $a(\kappa')$ agree.
2) For vertically adjacent $\kappa, \kappa' \in [k]^2$, the second components of $a(\kappa)$ and $a(\kappa')$ agree.
3) For all $\kappa \in \mathcal{C}$, we have $a(\kappa) \in \mathcal{T}(\kappa)$.

On the same inputs, we also define the problem $\#\mathrm{GridTiling}$, which asks to determine the *number* of grid tilings, and the problem $\oplus\mathrm{GridTiling}$, which asks to determine the *parity* of this number. All three problems are parameterized by k.

Remark II.5. Our definition of $\mathrm{GridTiling}$ is actually a generalized version of Marx's formulation: In his original definition, the set $\mathcal{C}$ is fixed to $\mathcal{C} = [k]^2$ on all instances, i.e., the third condition of Definition II.4 is required to apply for *all* $\kappa \in [k]^2$.

In the full version, we prove the following theorem, which serves as the main reduction source in the subsequent sections.

Theorem II.6. *The three variants of* $\mathrm{GridTiling}$ *are complete for* $\mathsf{W}[1]$, $\#\mathsf{W}[1]$ *or* $\oplus\mathsf{W}[1]$, *respectively. They admit no* $n^{o(k/\log k)}$ *time algorithms, even on instances with* $|\mathcal{C}| = \mathcal{O}(k)$, *unless* ETH, $\#\mathsf{ETH}$ *or* $\oplus\mathsf{ETH}$ *fails, respectively.*

We add an extension to Theorem II.6 that allows us to assume input instances to be balanced along rows or columns.

Lemma II.7. *Let* $\mathcal{A} = (n, k, \mathcal{C}, \mathcal{T})$ *be an instance for* $\mathrm{GridTiling}$ *and let* $\mathfrak{W}$ *be either of the words "horizontal" or "vertical". In polynomial time, we can then compute a number* $T \in \mathbb{N}$ *and an instance* $\mathcal{A}' = (n', k, \mathcal{C}, \mathcal{T}')$ *with* $n' = \mathcal{O}(k^2 n)$ *such that the instances* $\mathcal{A}$ *and* $\mathcal{A}'$ *have precisely the same grid tilings. For* $u \in [n]$, *write* $(u, \star) := \{(u, v) \mid v \in [n]\}$. *For* $v \in [n]$, *write* $(\star, v) := \{(u, v) \mid u \in [n]\}$. *Then we have the following:*

1) *If* $\mathfrak{W}$ *is "horizontal", then for all* $\kappa \in \mathcal{C}$ *and* $u \in [n']$, *we have* $|\mathcal{T}'(\kappa) \cap (u, \star)| = T$.
2) *If* $\mathfrak{W}$ *is "vertical", then for all* $\kappa \in \mathcal{C}$ *and* $v \in [n']$, *we have* $|\mathcal{T}'(\kappa) \cap (\star, v)| = T$.

III. Holants, matchgates, linear combinations of signatures

In the following, we give a introduction to what we call the *Holant framework*, a toolbox introduced by [53], [39], [54].

Definition III.1. A *signature graph* is an edge-weighted graph Ω which may feature parallel edges, and which has a *vertex function* $f_v : \{0,1\}^{I(v)} \to \mathbb{C}$ associated with each $v \in V(\Omega)$. We also call f_v the *signature* of v. If v has degree d and an edge-ordering $I(v) = \{e_1, \ldots, e_d\}$ is specified, we also consider $f_v : \{0,1\}^d \to \mathbb{C}$.

The *Holant* of Ω is a particular sum over edge assignments $x \in \{0,1\}^{E(\Omega)}$. For $x \in \{0,1\}^{E(\Omega)}$, we say that $e \in E(\Omega)$ is *active in* x if $x(e) = 1$ holds, and we tacitly identify x with the set of active edges in x. Given a subset $S \subseteq E(\Omega)$, we write $x|_S$ for the restriction of x to S, which is the unique assignment in $\{0,1\}^S$ that agrees with x on S.

Definition III.2 (adapted from [53])**.** Let Ω be a signature graph with edge weights $w : E(\Omega) \to \mathbb{C}$ and a vertex function $f_v : \{0,1\}^{I(v)} \to \mathbb{C}$ for each $v \in V(\Omega)$. For $x \in \{0,1\}^{E(\Omega)}$, we define

$$\mathrm{val}_\Omega(x) := \prod_{v \in V(\Omega)} f_v(x|_{I(v)}), \tag{1}$$

$$w_\Omega(x) := \prod_{e \in x} w(e), \tag{2}$$

and we say that x *satisfies* Ω if $\mathrm{val}_\Omega(x) \neq 0$ holds. Furthermore, we define

$$\mathrm{Holant}(\Omega) := \sum_{x \in \{0,1\}^{E(\Omega)}} w_\Omega(x) \cdot \mathrm{val}_\Omega(x). \tag{3}$$

A. Gates and matchgates

In some occasions, we can simulate signatures f appearing in a signature graph Ω by gadgets, i.e., signature graphs on "basic" signatures that realize f. We call such gadgets *gates*, similar to the $\mathcal{F}$-gates in [54], and we will be particularly interested in *matchgates*. These are gates Γ that feature, at each vertex $v \in V(\Gamma)$, the perfect matching signature that maps $x \in \{0,1\}^{I(v)}$ to

$$\mathtt{HW_{=1}}(x) := [\mathrm{hw}(x) = 1].$$

Definition III.3. A *gate* is a signature graph Γ containing a set $D \subseteq E(\Gamma)$ of *dangling edges*, all of which have edge-weight 1. A dangling edge is an edge e that is incident with only one vertex. We enumerate the dangling edges of Γ as $1, \ldots, |D|$.

Given a signature graph Ω, a vertex $v \in V(\Omega)$ of degree $|D|$, and a numbering of $I(v)$ as $I(v) = \{e_1, \ldots, e_{|D|}\}$, we can *insert* Γ at v by deleting v, placing a copy of Γ into G, and identifying e_i with the i-th dangling edge of Γ, for all i.

For disjoint sets A, B, and for $x \in \{0,1\}^A$ and $y \in \{0,1\}^B$, write $xy \in \{0,1\}^{A \cup B}$ for the assignment that agrees with x on A, and with y on B. We say that xy *extends* x. The *signature of* Γ is the function $\mathrm{Sig}(\Gamma) : \{0,1\}^D \to \mathbb{Q}$ that maps $x \in \{0,1\}^D$ to

$$\mathrm{Sig}(\Gamma, x) = \sum_{y \in \{0,1\}^{E(\Gamma) \setminus D}} w_\Gamma(xy) \cdot \mathrm{val}_\Gamma(xy). \tag{4}$$

We say that Γ *realizes* $\mathrm{Sig}(\Gamma)$. If all $v \in V(\Gamma)$ feature the function $\mathtt{HW_{=1}}$ defined above, then Γ is a *matchgate*. Finally, we call Γ planar if it can be drawn in the plane with all dangling edges on the outer face, such that they appear in the order $1, \ldots, |D|$ in a clockwise traversal of this face.

By the following lemma, if Γ realizes a signature f, and v is a vertex with signature f in a signature graph Ω, then we can insert Γ at v in a way that preserves Holants. In particular, we can use this to reduce $\mathrm{Holant}(\Omega)$ to $\mathrm{PerfMatch}$ if all signatures in Ω can be realized by matchgates.

Lemma III.4. *Let Ω be a signature graph, let $v \in V(\Omega)$ be arbitrary, and let f_v denote the vertex function of v in Ω. Furthermore, let Γ be a (match-)gate with $\mathrm{Sig}(\Gamma) = f_v$, and let Ω' be obtained from Ω by inserting Γ at v. Then* $\mathrm{Holant}(\Omega) = \mathrm{Holant}(\Omega')$. *If Ω and Γ are planar and Ω is given together with a plane embedding, then the following holds: If we order $I(v)$ according to its clockwise ordering in the embedding and insert Γ under this order, then Ω' is planar.*

We now consider specific matchgates that will be relevant later. To simplify, we abbreviate the following 4-bitstrings:

 := 0000, := 0101,  := 1010,  := 1111,
 := 1000, 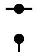 := 0010, 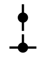 := 1101, := 0111.

In Figure 1, we define a signature $\mathtt{PASS}$ of arity 4 and two signatures $\mathtt{PRE}$ and $\mathtt{ACT}$ of arity 6. Note that $\mathtt{PASS}$ essentially acts as a "crossing" signature: It enforces equality on its western and eastern dangling edges (numbered 4 and 2), as well as on its northern and southern dangling edges (numbered 1 and 3). However, if all dangling edges are active, then the output of $\mathtt{PASS}$ is -1 rather than 1. This flipped sign allows $\mathtt{PASS}$ to admit a planar matchgate $\Gamma_{\mathtt{PASS}}$, shown in Figure 1. We verified that $\mathrm{Sig}(\Gamma_{\mathtt{PASS}}) = \mathtt{PASS}$ holds by means of a computer program: For all $x \in \{0,1\}^4$, we showed mechanically that $\mathrm{Sig}(\Gamma_{\mathtt{PASS}}, x) = \mathtt{PASS}(x)$ holds. Note that this verification can also be carried out by hand. It should also be noted that planar matchgates for $\mathtt{PASS}$ were already studied in [53], [41].

Next, we consider the signatures $\mathtt{PRE}$ and $\mathtt{ACT}$, each of arity 6. We consider their last two inputs (the dangling edges with

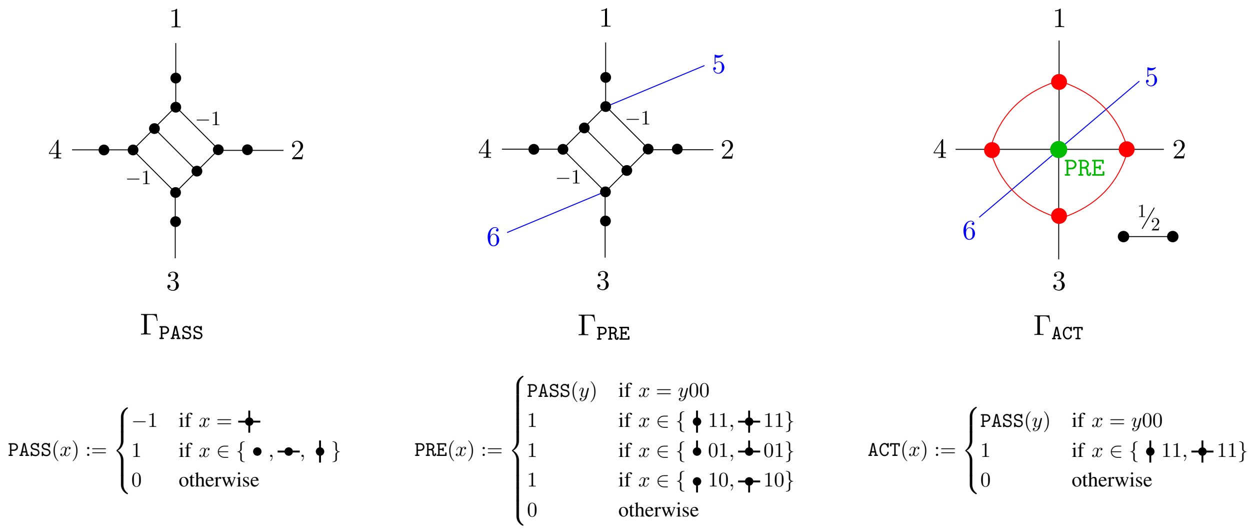

Fig. 1. The matchgates $\Gamma_{\texttt{PASS}}$, $\Gamma_{\texttt{PRE}}$ and $\Gamma_{\texttt{ACT}}$ and the signatures $\texttt{PASS}$, $\texttt{PRE}$ and $\texttt{ACT}$. Note that $\Gamma_{\texttt{PASS}}$ has four dangling edges, numbered 1 to 4, whereas $\Gamma_{\texttt{PRE}}$ and $\Gamma_{\texttt{ACT}}$ each have six dangling edges, numbered 1 to 6. The signature $\texttt{PASS}$ is defined on assignments $x \in \{0,1\}^4$, while $\texttt{PRE}$ and $\texttt{ACT}$ are defined on assignments $x \in \{0,1\}^6$. These strings correspond canonically to assignments at the dangling edges of $\Gamma_{\texttt{PASS}}$, $\Gamma_{\texttt{PRE}}$ and $\Gamma_{\texttt{ACT}}$. All black vertices are assigned $\texttt{HW}_{=1}$. In the gate $\Gamma_{\texttt{ACT}}$, all red vertices are assigned $\texttt{PASS}$, and the green middle vertex is assigned $\texttt{PRE}$. Note that we can also view $\Gamma_{\texttt{ACT}}$ as a matchgate by realizing its signatures with the matchgates $\Gamma_{\texttt{PASS}}$ and $\Gamma_{\texttt{ACT}}$. All matchgates are planar after removal of the dangling edges 5 and 6, which will later connect to apex vertices.

numbers 5 and 6) as "switches", which will later be connected to apices. It is crucial to observe that

$$\texttt{PRE}(x00) = \texttt{ACT}(x00) = \texttt{PASS}(x) \quad \forall x \in \{0,1\}^4.$$

That is, if the two switch edges are not active, then $\texttt{PRE}$ and $\texttt{ACT}$ behave exactly like $\texttt{PASS}$ on their non-switch inputs. If both switches are active, then some differences occur, namely, the restriction to non-switch edges must be in state ϕ or ┼ for $\texttt{PRE}$ or $\texttt{ACT}$ to yield a nonzero value. Furthermore, if only one of the two switches is active, then $\texttt{ACT}$ yields value zero, while $\texttt{PRE}$ still allows such assignments (such as ┴01). We verified with a computer program that $\texttt{PRE} = \mathrm{Sig}(\Gamma_{\texttt{PRE}})$ holds for the matchgate $\Gamma_{\texttt{PRE}}$ from Figure 1. In the full version, we prove manually that $\texttt{ACT} = \mathrm{Sig}(\Gamma_{\texttt{ACT}})$ holds.

B. Linear combinations of matchgate signatures

We introduce our main tool for the later sections, a technique that allows us to simulate signatures by linear combinations of other signatures, in particular, of matchgate signatures.

Definition III.5. Let $f = c_1 \cdot f_1 + \ldots + c_t \cdot f_t$ be a signature, where $c_1, \ldots, c_t \in \mathbb{C}$ are coefficients and $f_1, \ldots, f_t$ are signatures, and the linear combination is point-wise. Then we say that f is t-combined from constituents $f_1, \ldots, f_t$.

We apply such linear combinations as follows: Assume we are given a signature graph that features k occurrences of some interesting signature f which cannot be realized by matchgates. If we can express f as a linear combination of t constituents that do admit matchgates, then the following lemma allows us to compute $\mathrm{Holant}(\Omega)$ from the Holants of t^k derived signature graphs whose signatures all admit matchgates.

Lemma III.6. *Let Ω be a signature graph, let $k, t \in \mathbb{N}$ and let $w_1, \ldots, w_k$ be distinct vertices of Ω such that the following holds: For all $\kappa \in [k]$, the signature f_κ at w_κ admits coefficients $c_{\kappa,1}, \ldots, c_{\kappa,t} \in \mathbb{C}$ and signatures $g_{\kappa,1}, \ldots, g_{\kappa,t}$ such that $f_\kappa = \sum_{i=1}^{t} c_{\kappa,i} \cdot g_{\kappa,i}$. Given a tuple $\theta \in [t]^k$, let Ω_θ be defined by replacing, for each $\kappa \in [k]$, the vertex function f_κ at w_κ with $g_{\kappa,\theta(\kappa)}$. Then we have*

$$\mathrm{Holant}(\Omega) = \sum_{\theta \in [t]^k} \left(\prod_{\kappa=1}^{k} c_{\kappa,\theta(\kappa)} \right) \cdot \mathrm{Holant}(\Omega_\theta). \tag{5}$$

When using Lemma III.6 for positive results, as in Section IV, then the right-hand side of (5) is "easy", in the sense that the values $\mathrm{Holant}(\Omega_\theta)$ for all θ can be obtained efficiently, e.g., by reduction to planar $\mathrm{PerfMatch}$. In the same way, Lemma III.6 also allows us to prove hardness results under Turing reductions, as we do in Sections V and VI: In this case, the left-hand side is "hard" and could be computed from oracle access to the values $\mathrm{Holant}(\Omega_\theta)$ for all θ.

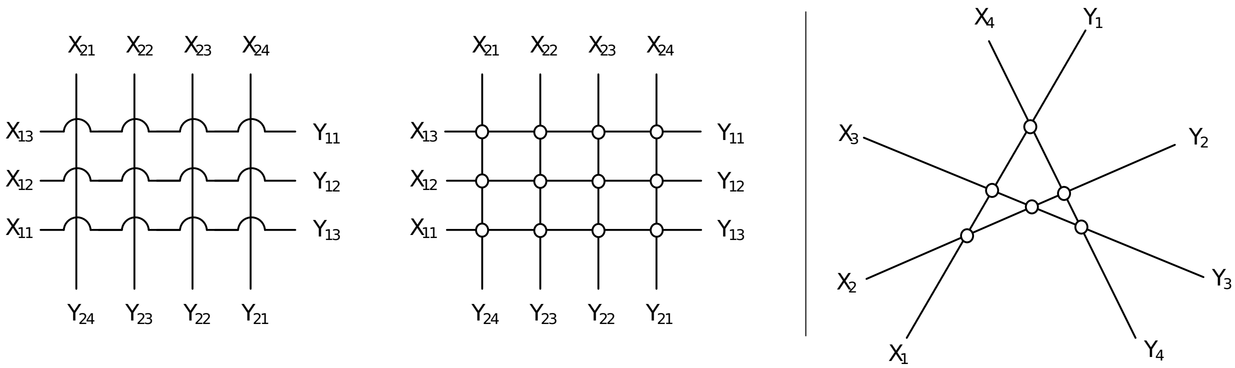

Fig. 2. The first two subfigures show a grid cap and the matchgate realizing one of the constituents used to realize the grid cap. The third subfigure shows the matchgate used to simulate a cross cap. In these matchgates, all vertices are assigned the signature PASS.

IV. PerfMatch on bounded-genus graphs

In this section, we present a first application of the framework of combined signatures: We show that, for graphs of genus k, the quantity $\mathrm{PerfMatch}(G)$ can be expressed as a linear combination of 4^k values $\mathrm{PerfMatch}(G_i)$, where G_i is a planar graph for all $i \in [4^k]$. The linear combinations resemble those used in [12], [13], [14], but unlike these papers, we can state our linear combinations without any necessity for Pfaffian orientations.

Following [13], we assume that the graph G in question is given to us together with a plane model: All vertices of G are drawn in a polygon P with $2k$ sides. If there is a bunch of d_i parallel edges $x_i = x_{i1}x_{i2}\cdots x_{id_i}$ leaving P from one side and going into P through another side, we denote the two sides by a_i and a_i^{-1} respectively. Since the edges are parallel, when we walk along the sides of P counterclockwise, we meet the exits of edges in the order $x_{i1}x_{i2}\cdots x_{id_i}$ on side a_i, then the entrances of edges in the order $x_{id_i}x_{i(d_i-1)}\cdots x_{i1}$ on side a_i^{-1}.

If G can be embedded on an orientable compact boundaryless surface S of genus k, then it can be drawn such that there are no edges crossing inside P, and the sides of P are $a_1a_2a_1^{-1}a_2^{-1}a_3a_4a_3^{-1}a_4^{-1}\cdots a_{2k-1}a_{2k}a_{2k-1}^{-1}a_{2k}^{-1}$. The side pair a_i, a_i^{-1} represents boundaries to be glued together. When G is drawn on the surface S, the edge bunches x_1 and x_2 overpass each other without any edges crossing; see the left picture of Figure 2 for such a situation, which we call a *grid cap*.

We use linear combinations of matchgates (like the one shown in the middle of Figure 2) to simulate the grid cap by a planar graph. Write x_i^{-1} to denote $x_{id_i}x_{i(d_i-1)}\cdots x_{i1}$. Then the grid cap realizes a function that is defined on assignments (x_1, x_2, y_1, y_2) to its dangling edges as follows:

$$O(x_1, x_2, y_1, y_2) = [y_1 = x_1^{-1}] \cdot [y_2 = x_2^{-1}].$$

The straightforward idea is to place a PASS matchgate at each crossing of overpassing edges, as shown in the middle of Figure 2. Let us denote by $C(x_1, x_2, y_1, y_2)$ the signature of the resulting gate. In any satisfying assignment (x_1, x_2, y_1, y_2) to its dangling edges, there are $\mathrm{hw}(x_1) \cdot \mathrm{hw}(x_2)$ instances of PASS in state ╋, each of which gives a factor -1, while all other instances of PASS (in states ┃, ━, •) give a factor 1, so

$$C(x_1, x_2, y_1, y_2) = (-1)^{\mathtt{ODD}(x_1)\cdot\mathtt{ODD}(x_2)} \cdot [y_1 = x_1^{-1}] \cdot [y_2 = x_2^{-1}].$$

Lemma IV.1. *Every grid cap gate is a linear combination of* 4 *matchgates, given by*

$$O(x_1, x_2, y_1, y_2) = \frac{1}{2}(1 + (-1)^{\mathtt{ODD}(x_1)} + (-1)^{\mathtt{ODD}(x_2)} + (-1)^{\mathtt{ODD}(x_1)+\mathtt{ODD}(x_2)+1}) \cdot C(x_1, x_2, y_1, y_2).$$

Note that we can indeed realize the four constituent signatures via planar matchgates: To simulate, e.g., the product of $(-1)^{\mathtt{ODD}(x_1)}$ with $C(x_1, x_2, y_1, y_2)$, it suffices to assign edge-weight -1 to one horizontal edge in each row of the matchgate realizing C.

We now turn our attention to non-orientable surfaces and their plane models: If G can be embedded on a non-orientable surface S, which is the connected sum of a surface of orientable genus k with either a projective plane or a Klein bottle, then it can be drawn without crossings inside P, such that the sides of P are $a_1a_2a_1^{-1}a_2^{-1}a_3a_4a_3^{-1}a_4^{-1}\cdots a_{2k-1}a_{2k}a_{2k-1}^{-1}a_{2k}^{-1}a_{2k+1}a_{2k+2}$ and $a_1a_2a_1^{-1}a_2^{-1}a_3a_4a_3^{-1}a_4^{-1}\cdots a_{2k-1}a_{2k}a_{2k-1}^{-1}a_{2k}^{-1}a_{2k+1}a_{2k+2}\ a_{2k+3}a_{2k+4}$ respectively. Here, the side pair a_ia_i means that, when a bunch of edges $x_i = x_{i1}x_{i2}\cdots x_{id_i}$ leaves the interior of P through the first side a_i and then enters back into P through the second side a_i, then we meet the exits and entrances in the order x_ix_i. Such a bunch of edges is called a *cross* *cap*, and it realizes a function

$$O(x,y)=[y=x].$$

If we draw it on the plane and replace each crossing by a PASS matchgate, as shown in the right part of Figure 2, we get a matchgate realizing

$$C(x,y)=(-1)^{\binom{\mathrm{hw}(x)}{2}}\cdot[y=x].$$

Lemma IV.2. *Every cross cap gate is a linear combination of 2 matchgates, given by*

$$O(x,y)=\frac{1-i}{2}\cdot i^{\mathrm{hw}(x)}\cdot C(x,y)+\frac{1+i}{2}\cdot(-i)^{\mathrm{hw}(x)}\cdot C(x,y).$$

Using the fact that G is embedded as a plane model, and using the combined signatures for grid caps and cross caps from the last two lemmas, we then obtain the following known theorem.

Theorem IV.3. *[13] Let G be a graph that is embedded on a surface. Then $\mathrm{PerfMatch}(G)$ is a summation of $\mathrm{PerfMatch}$ of 2^{2k}, 2^{2k+1} or 2^{2k+2} planar graphs, respectively, if the surface is the connected sum of an orientable surface of genus k with the plane, the projective plane, or the Klein bottle, respectively.*

Proof. By Lemma IV.1 and IV.2, use Lemma III.6 on the k grid caps and 0, 1 or 2 cross caps. □

The essence of Lemma IV.1 is that we can use the four matchgates to realize all four columns of the basis

$$\begin{pmatrix}1 & 1\\ 1 & -1\end{pmatrix}^{\otimes 2},$$

so that we can then obtain any other function by linear combinations. The same observation also holds for a larger base

$$\begin{pmatrix}1 & 1\\ 1 & -1\end{pmatrix}^{\otimes m}.$$

We give an example: In a cross cap of m edges, we may replace each edge by a bunch of parallel edges, and call the result a *grated cross cap*. All the $\binom{m}{2}$ latent crossings of the cross cap become grid caps in the grated cross cap.

Fact IV.4. *Every grated cross cap gate over m bunches of edges, as defined above, can be expressed as a linear combination of 2^m planar matchgates.*

In fact, these 2^m basis matchgates are powerful enough to express (as a linear combination) any function that depends only upon the parities $p_1,\dots,p_m$ of active edges in the m edge bunches. However, among these functions, we currently only know one interesting function, i.e., the grid cap. Even the grated cross cap seems too artificial to be related with a natural tractability result. A similar generalization applies to Lemma IV.2, where the functions to be expressed may also depend upon residuals of the numbers of active edges in the m edge bunches, in this case however modulo 4 rather than 2.

V. The permanent on k-apex graphs

In this section, we prove Theorem I.1 by using combined signatures for a reduction from $\#\mathrm{GridTiling}$ to the permanent on k-apex graphs: First, in Section V-A, we express an instance of $\#\mathrm{GridTiling}$ as $\mathrm{Holant}(G)$ for a signature graph G. Then we realize the signatures of G as combined signatures in Section V-B. Parts of this section will be reused in Section VI with an added layer of technicalities.

A. Global construction

In the following, let $\mathcal{A}=(n,k,\mathcal{C},\mathcal{T})$ be a fixed instance to $\#\mathrm{GridTiling}$, as specified in Definition II.4. By applying vertical balance (see Lemma II.7), we may assume the existence of some number $T\le n$ such that for all $\kappa\in\mathcal{C}$ and all $v\in[n]$, there are exactly T elements of type $(\star,v)$ in $\mathcal{T}(\kappa)$. This will become relevant in Section V-B.

First, we reformulate $\mathcal{A}$ as the Holant of a signature graph $G=G(\mathcal{A})$. This graph G consists of a $k\times k$ square grid of *cells,* and $4k$ additional *border vertices* adjacent to the borders of the grid, as seen in the left part of Figure 3. Note that G is planar. We denote its vertices by c_κ for $\kappa\in\Xi$, where $\Xi:=[k]^2\cup\{\mathsf{N},\mathsf{W},\mathsf{S},\mathsf{E}\}\times[k]$.

For $i\in[k]$, we declare (N,i) to be vertically adjacent to $(1,i)$, and (S,i) to (k,i). Likewise, we declare (W,i) to be horizontally adjacent to $(i,1)$, and (E,i) to (i,k). We refer to the neighbors of any index $\kappa\in\Xi$ or vertex $c_\kappa\in V(G)$ using cardinal directions in the obvious way, e.g., we may speak of the northern neighbor of a vertex. Between any pair of vertices c_κ and $c_{\kappa'}$ with adjacent indices κ and κ', we place a set $E_{\kappa,\kappa'}$ of n parallel edges, which we call an *edge bundle*.

We proceed to define the signatures of G. In the assignments $a\in\{0,1\}^{E(G)}$ that we are interested in, each edge bundle features exactly one active edge, which is used to encode a number from $[n]$. At border vertices, we place the signature $\mathtt{HW}_{=1}$ to ensure this. The signatures of cells c_κ with $\kappa\in[k]^2$ are then defined so that each cell propagates the number $x_W\in[n]$

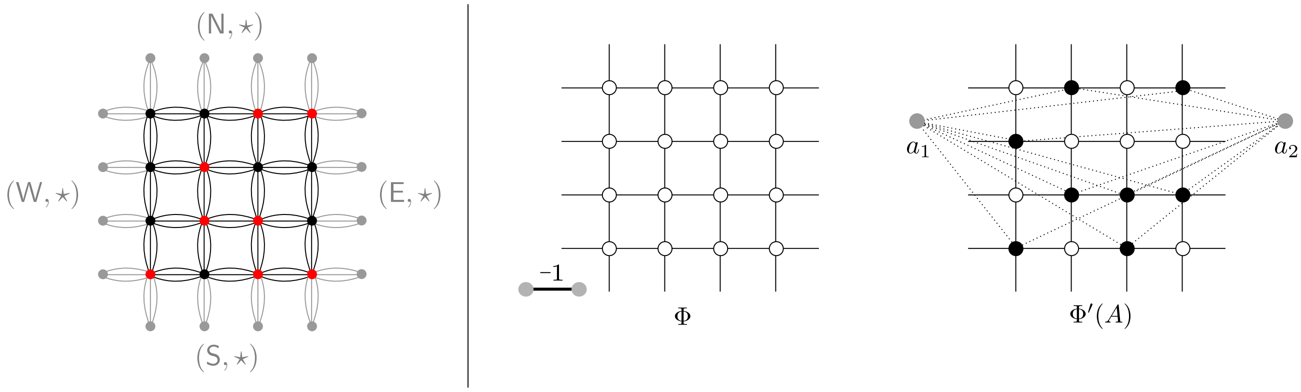

Fig. 3. The left part of the figure shows the signature graph $G(\mathcal{A})$. Border vertices c_κ for $\kappa \in \{\mathsf{N}, \mathsf{W}, \mathsf{S}, \mathsf{E}\} \times [k]$ and their incident edges are colored gray. Cell vertices c_κ for $\kappa \in \mathcal{C}$ are colored red, while vertices c_κ for $\kappa \in [k]^2 \setminus \mathcal{C}$ are colored black. Horizontally or vertically adjacent vertices are connected by an edge bundle of n parallel edges. The right part of the figure shows the gates Φ and $\Phi'(A)$. Each white vertex is assigned PASS, each black vertex is assigned ACT, and each gray vertex is assigned $\mathtt{HW}_{=1}$. Edges from apices in Φ' are drawn dashed to avoid visual cluttering. Note that, due to the balance property of $\mathcal{T}$, we may assume that every column has the same number T of occurrences of ACT.

encoded by its western incident edge bundle to the east, and its northern number $x_N \in [n]$ to the south, while checking along the way whether $(x_W, x_N) \in \mathcal{T}(\kappa)$ holds.

Remark V.1. We adhere to the following conventions in this section:

- For $v \in [n]$, we often identify the string $0^{v-1}10^{n-v} \in \{0,1\}^n$ with the number v when it is clear from the context which of these two objects we currently refer to.
- For $\kappa \in [k]^2$, the $4n$ incident edges of each vertex c_κ are ordered such that all northern edges appear first, in a block of length n, followed by the n eastern, the n southern, and finally the n western edges.
- We implicitly consider strings $x \in \{0,1\}^{4n}$ to be decomposed into $x = x_N x_E x_S x_W$ with $x_N, x_E, x_S, x_W \in \{0,1\}^n$.

Using these conventions, we then define the following predicates for strings $x \in \{0,1\}^{4n}$:

$$\begin{aligned}\varphi_{one}(x) &\equiv \mathrm{hw}(x_N) = 1 \wedge \mathrm{hw}(x_W) = 1,\\ \varphi_{prop}(x) &\equiv x_N = x_S \wedge x_W = x_E.\end{aligned}$$

If a function f satisfies $\varphi_{prop}(x)$ for each $x \in \mathrm{supp}(f)$, then we call f *propagating*. For each $\kappa \in [k]^2$, we place a specific propagating signature f_κ at the vertex c_κ in order to complete G to a signature graph whose satisfying assignments correspond bijectively to the grid tilings of $\mathcal{A}$.

Definition V.2. For all $\kappa \in [k]^2 \setminus \mathcal{C}$, we define the vertex function $f_\kappa : \{0,1\}^{4n} \to \{0,1\}$ of c_κ such that, for all $x \in \{0,1\}^{4n}$ satisfying the predicate $\varphi_{one}(x)$, we have

$$f_\kappa(x) := [\varphi_{prop}(x)].$$

Note that no requirement is imposed upon $f_\kappa(x)$ on those $x \in \{0,1\}^{4n}$ that fail to satisfy $\varphi_{one}(x)$. For all remaining κ, namely all $\kappa \in \mathcal{C}$, we define the vertex function g_κ of c_κ on such $x \in \{0,1\}^{4n}$ by declaring

$$g_\kappa(x) := [\varphi_{prop}(x) \wedge (x_W, x_N) \in \mathcal{T}(\kappa)]$$

In the following, we show that $G = G(\mathcal{A})$ indeed encodes $\mathcal{A}$ properly.

Lemma V.3. *The grid tilings of $\mathcal{A}$ correspond bijectively to the satisfying assignments $x \in \{0,1\}^{E(G)}$ of G, and each satisfying assignment x additionally has $\mathrm{val}_G(x) = 1$.*

In the next subsection, we realize each signature f_κ for $\kappa \in \mathcal{C}$ as a planar matchgate, and each g_κ for $\kappa \in [k]^2 \setminus \mathcal{C}$ as a linear combination of two matchgate signatures that have maximum apex number 2. Note that the remaining signatures $\mathtt{HW}_{=1}$ occurring in G are planar.

B. Realizing cell signatures

It can be shown (under no additional assumptions) that some of the signatures g_κ for $\kappa \in [k]^2$ are non-planar. From a complexity viewpoint, if all such signatures were planar and we knew explicit planar matchgates, then we could reduce #GridTiling to planar PerfMatch, and thus show $\mathsf{FP} = \#\mathsf{P}$ by the FKT method. Rather than trying to use planar matchgates, we show that each signature g_κ can be realized as a specific *linear combination* of the signatures of one planar and one 2-apex matchgate. (At least one non-planar constituent is necessary, as we could otherwise show $\mathsf{FPT} = \#\mathsf{W}[1]$.)

In the remainder of this section, we consider $\kappa \in [k]^2$ to be fixed, we write $A = \mathcal{T}(\kappa)$ and we recall that $A \subseteq [n]^2$. The constituents for g_κ will be the signatures of two gates Φ and $\Phi'(A)$, which use as building blocks the signatures PASS and ACT from Section III.

Definition V.4. Let $n \in \mathbb{N}$ and let $A \subseteq [n]^2$. We define gates Φ and $\Phi' = \Phi'(A)$ with $4n$ dangling edges (that is, with n dangling edges for each cardinal direction) as follows. Consider also the right part of Figure 3.

- To obtain the gate Φ, arrange vertices b_τ for $\tau \in [n]^2$ in a $n \times n$ grid and assign the signature PASS to each such vertex. Add a single edge of weight -1 between two fresh vertices of signature $\mathtt{HW}_{=1}$.
- A similar construction yields the gate Φ': Starting from Φ, remove the extra edge of weight -1, add apex vertices a_1 and a_2 with signatures $\mathtt{HW}_{=1}$, and for all $\tau \in A$, do the following:
  1) Replace the signature PASS at b_τ with ACT.
  2) Add the edges $a_1 b_\tau$ and $a_2 b_\tau$ and declare these to be the last two edges in the edge ordering of $I(v_\tau)$.

Recall that PASS is realized by the planar matchgate $\Gamma_{\mathtt{PASS}}$, so we can also view the gate Φ as a planar matchgate after realizing all signatures by matchgates. We will later switch between these views depending on the application. Note also that the 2-coloring of $\Gamma_{\mathtt{PASS}}$ can be extended to one of Φ. Likewise, ACT is realized by the matchgate $\Gamma_{\mathtt{ACT}}$, which is planar when ignoring its dangling edges 5 and 6. That is, after realizing each occurrence of ACT by $\Gamma_{\mathtt{ACT}}$, the resulting matchgate obtained from Φ' is planar after removal of a_1 and a_2. Furthermore, a_1 is only adjacent to green-colored vertices of $\Gamma_{\mathtt{ACT}}$, while a_2 is only adjacent to red-colored vertices of $\Gamma_{\mathtt{ACT}}$, so Φ' admits a valid 2-coloring.

Our goal for this subsection is to realize the signatures f_κ and g_κ from Definition V.2. In the following, we prove that $f_\kappa = \mathrm{Sig}(\Phi)$ and that g_κ can be realized by a linear combination of $\mathrm{Sig}(\Phi)$ and $\mathrm{Sig}(\Phi')$. It will be crucial for our calculations to assume our instance $\mathcal{A}$ for GridTiling to be balanced: By Lemma II.7, we assume there is some $T \in \mathbb{N}$ such that $|A \cap (\star, v)| = T$ for all $v \in [n]$. That is, in the right part of Figure 3, we may assume that every column of $\Phi'(A)$ features the same number T of vertices with signature ACT.

Lemma V.5. *Recall the definition of the predicates φ_{one} and φ_{prop} on the preceding page. Let $x \in \{0,1\}^{4n}$ be an assignment that satisfies the predicate φ_{one}. Then*

$$\mathrm{Sig}(\Phi, x) = \begin{cases} 0 & \text{if } \neg\varphi_{prop}(x), \\ 1 & \text{if } \varphi_{prop}(x). \end{cases} \tag{6}$$

$$\mathrm{Sig}(\Phi'(A), x) = \begin{cases} 0 & \text{if } \neg\varphi_{prop}(x) \\ \begin{cases} -T & \text{if } (x_W, x_N) \notin A \\ -T+2 & \text{if } (x_W, x_N) \in A \end{cases} & \text{if } \varphi_{prop}(x). \end{cases} \tag{7}$$

Note that $f_\kappa = \mathrm{Sig}(\Phi)$ for $\kappa \in [k]^2 \setminus \mathcal{C}$. For $\kappa \in \mathcal{C}$ and for $x \in \{0,1\}^{4n}$ satisfying φ_{one}, we have

$$g_\kappa(x) = \frac{T}{2} \cdot \mathrm{Sig}(\Phi, x) + \frac{1}{2} \cdot \mathrm{Sig}(\Phi'(\mathcal{T}(\kappa)), x). \tag{8}$$

Let us show how Lemma V.5 implies Theorem I.2. We will require parts of this argument again in Section VI.

Proof of Theorem I.2. By Lemma V.3, we know that $\mathrm{Holant}(G)$ counts the grid tilings of $\mathcal{A}$. Using the linear combination (8) and Lemma III.6 about the linear combinations of signatures, we obtain

$$\mathrm{Holant}(G) = \frac{1}{2^{|\mathcal{C}|}} \sum_{\omega : \mathcal{C} \to [2]} T^{d(\omega)} \cdot \mathrm{perm}(H_\omega). \tag{9}$$

For $\omega : \mathcal{C} \to [2]$, the number $d(\omega)$ is the number of 1-entries in ω, and the graph H_ω is obtained as follows:

- For $\kappa \in [k]^2 \setminus \mathcal{C}$ and for $\kappa \in \mathcal{C}$ with $\omega(\kappa) = 1$, insert the matchgate Φ at the cell vertex c_κ.
- For $\kappa \in \mathcal{C}$ with $\omega(\kappa) = 2$, insert the matchgate $\Phi'(\mathcal{T}(\kappa))$ at c_κ.

Since G is planar, and since Φ is planar and $\Phi'(\mathcal{T}(\kappa))$ for $\kappa \in \mathcal{C}$ has at most 2 apices, it follows that $\mathrm{apex}(H_\omega) \le 2|\mathcal{C}|$ for all $\omega : \mathcal{C} \to [2]$, and this proves the required parameter bound. By 2-coloring the matchgates Φ and Φ', it can be verified that each graph H_ω is bipartite. Additionally, by construction of the matchgates $\Gamma_{\mathtt{PASS}}$ and $\Gamma_{\mathtt{ACT}}$, every graph H_ω features only edge-weights from the set $\{-1, \frac{1}{2}, 1\}$. Non-unit edge-weights in H_ω appear only at edges $uv \in E(H_\omega)$ not incident with apices. We can hence use standard weight simulation techniques to remove the edge-weights -1 and $\frac{1}{2}$, as in [1], while maintaining the apex number. We consequently obtain $\#\mathsf{W}[1]$-completeness of the permanent under the apex parameter and the claimed lower bound under $\#\mathsf{ETH}$. ☐

Remark V.6. The following might prove useful for later applications: By construction, the apices in the constructed graphs H_ω form an independent set, for any $\omega : [k]^2 \to [2]$, and each non-apex vertex in H_ω is incident with at most one apex. This last condition holds because the matchgate $\Gamma_{\mathtt{ACT}}$ has no vertex with two incident dangling edges.

VI. The permanent modulo 2^k

We prove Theorem I.3, which asserts $\oplus$W[1]-hardness of evaluating the permanent mod 2^k. We reduce from the problem $\oplus\mathrm{GridTiling}$, the parity version of $\mathrm{GridTiling}$ from Definition II.4. From a high level, the proof resembles that of Theorem I.2, but the setting of modular evaluation requires us to apply linearly combined signatures in a more intricate way.

A. The main idea

Our reduction is based upon the following observation: Let $\mathcal{A} = (n, k, \mathcal{C}, \mathcal{T})$ be an instance for $\oplus\mathrm{GridTiling}$. For $\omega : \mathcal{C} \to [2]$, recall the graphs H_ω and the numbers $d(\omega)$ from the last section. We can rewrite (9) as

$$2^{|\mathcal{C}|} \cdot \#\mathrm{GridTiling}(\mathcal{T}) = \sum_{\omega:\mathcal{C}\to[2]} T^{d(\omega)} \cdot \mathrm{perm}(H_\omega). \tag{10}$$

Theorem II.6 asserts that computing $\oplus\mathrm{GridTiling}(\mathcal{T})$ is $\oplus$W[1]-hard. Let $M := 2^{|\mathcal{C}|}$ and assume we could evaluate $\mathrm{perm}(H_\omega)$ modulo $2M$ for all ω. Using arithmetic in $\mathbb{Z}/2M\mathbb{Z}$, we could then evaluate the entire right-hand-side of (10), and this allows us to compute $M \cdot \#\mathrm{GridTiling}(\mathcal{T})$ modulo $2M$, which is 0 iff $\#\mathrm{GridTiling}(\mathcal{T})$ is even, and M iff it is odd. Hence, it seems that we could solve $\oplus\mathrm{GridTiling}(\mathcal{T})$ with an oracle for the permanent modulo $2M = 2^{|\mathcal{C}|+1}$, and we might be tempted to believe that we just proved Theorem I.3.

However, the above argument suffers from a fatal gap: The graphs H_ω from the previous section feature edges of weight $\frac{1}{2}$, a number that does not exist in the rings $\mathbb{Z}/2^k\mathbb{Z}$ for $k \in \mathbb{N}$. In other words, the proof fails for the surprisingly philosophical reason that the instances H_ω constructed in the previous section do not even *exist* modulo 2^k. More precisely, the matchgate $\Gamma_{\mathtt{ACT}}$ used to realize the signature ACT features this offending weight. To obtain graphs H_ω that avoid edge-weights with even denominators, we therefore construct cell gates using the signature PRE rather than its more benign version ACT. This adds several complications to our arguments, which however vanish after a suitable linear combination.

B. Revisiting the cell gate

Let $A \subseteq [n]^2$ be fixed in the following, and recall the gates Φ and Φ' from Definition V.4. Note that Φ features only occurrences of PASS, which is realized by the matchgate $\Gamma_{\mathtt{PASS}}$ on edge-weights -1 and 1. We can therefore also realize this gate modulo 2^k. This does not apply to the gate $\Phi'(A)$, as the matchgate $\Gamma_{\mathtt{ACT}}$ realizing ACT features the weight $\frac{1}{2}$. We modify $\Phi'(A)$ to a new gate $\Gamma(A)$ by replacing all occurrences of ACT with PRE.

Definition VI.1. For $A \subseteq [n]^2$, let the gate $\Gamma(A)$ on $4n$ dangling edges be defined exactly as the gate $\Phi'(A)$ from Definition V.4, but replace every occurrence of ACT by PRE. For all $u, v \in [n]$, let $\alpha_{u,v}$ denote the number of occurrences of PRE among vertices b_τ with $\tau \in \{(1, v), \ldots, (u-1, v)\}$. Likewise, let $\beta_{u,v}$ denote the number of occurrences of PRE among vertices b_τ with $\tau \in \{(u+1, v), \ldots, (n, v)\}$.

Figuratively speaking, $\alpha_{u,v}$ is the number of occurrences of PRE in the column above (u, v), and $\beta_{u,v}$ is the number of occurrences below it. In Section V-B, we used the vertical balance property to ensure that $\alpha_{u,v} + \beta_{u,v}$ is equal to $T - 1$ when $(u, v) \in A$, and equal to T when $(u, v) \notin A$. In this section, this vertical balance will not be required, but *horizontal* balance will prove useful instead, for different reasons. For the remainder of our proofs, we define the following auxiliary polynomials, for all $u, v, w \in [n]$:

$$q_u := \sum_{z\in[n]} \alpha_{u,z} \cdot \beta_{u,z} - \binom{\alpha_{u,z}}{2} - \binom{\beta_{u,z}}{2}, \tag{11}$$

$$p_{u,v,w} := (\alpha_{u,v} - \beta_{u,v}) \cdot (\beta_{u,w} - \alpha_{u,w}), \tag{12}$$

$$r_{u,v} := \sum_{\substack{z\in[n]\setminus\{v\} \\ (u,z)\in A}} \beta_{u,z} + \alpha_{u,z}, \tag{13}$$

Using these polynomials, we can express the signature of Γ.

Lemma VI.2. *Let $A \subseteq [n]^2$, let $\Gamma = \Gamma(A)$ and let $x \in \{0,1\}^{4n}$ satisfy φ_{one}. Recall the conventions from Remark V.1, including that we implicitly decompose the string x into x_N, x_E, x_S, x_W.*

- *If $x_W \neq x_E$ or* $\mathrm{hw}(x_S) \neq 1$*, then* $\mathrm{Sig}(\Gamma, x) = 0$*.*

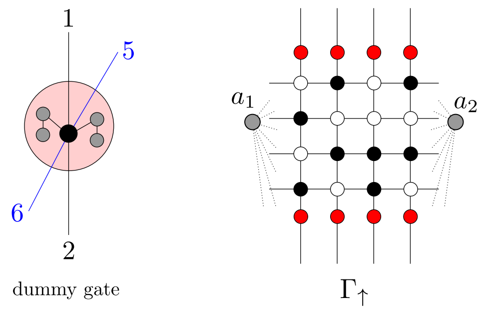

Fig. 4. A dummy gate is shown on the left. On the right, we see $\Gamma_\uparrow$, which is obtained from Γ by adding rows of dummy gates, shown red. Each gray vertex is assigned $\mathtt{HW}_{=1}$, and the apices connect to all black vertices (assigned PRE) and all red vertices (whose signature is realized by the dummy gate). White vertices are assigned PASS, and they are not adjacent to apices.

- *If $\varphi_{prop}(x)$ is true (i.e., we have $x_W = x_E$ and additionally $x_N = x_S$), write $u := x_W$ and $v := x_N$, with $u, v \in [n]$. Note that these numbers are well-defined. We call such assignments x* wanted, *and we have*

$$\mathrm{Sig}(\Gamma, x) = \begin{cases} q_u - r_{u,v} - \alpha_{u,v} - \beta_{u,v} & \text{if } (u,v) \notin A \\ q_u - r_{u,v} + 1 & \text{if } (u,v) \in A \end{cases}$$

- *If $\varphi_{prop}(x)$ is false (i.e, we have $x_W = x_E$, but $x_N \neq x_S$), then write $u := x_W$, $v := x_N$, and $w := x_S$. We call such assignments x* unwanted, *and we have*

$$\mathrm{Sig}(\Gamma, x) = \begin{cases} p_{u,v,w} & \text{if } (u,v) \notin A,\ (u,w) \notin A \\ p_{u,v,w} + \alpha_{u,v} - \beta_{u,v} & \text{if } (u,v) \notin A,\ (u,w) \in A \\ p_{u,v,w} + \beta_{u,w} - \alpha_{u,w} & \text{if } (u,v) \in A,\ (u,w) \notin A \\ p_{u,v,w} + \beta_{u,w} - \alpha_{u,w} + \alpha_{u,v} - \beta_{u,v} + 1 & \text{if } (u,v) \in A,\ (u,w) \in A \end{cases}$$

We note that the gate Γ essentially discriminates between six different assignment types, depending on whether x is wanted (giving 2 types) or unwanted (giving 4 types, depending on whether (x_W, x_N) and (x_W, x_S) are each contained in A). However, the actual value of $\mathrm{Sig}(\Gamma, x)$ is *not constant* for each of the six types, as it depends on u, v, w and the concrete values for $\alpha_{u',v'}$ and $\beta_{u',v'}$ for all $u', v' \in [n]$. Compare this to the gate Φ' from the previous section, which attains one of the three fixed values $\{0, -T, -T+2\}$. The remainder of the proof therefore aims at the following two goals:

Goal 1: Ensure that unwanted cases cancel out

Goal 2: Ensure that wanted cases do not depend upon the actual value of (x_W, x_N), but only on the information whether $(x_W, x_N) \in A$ or $(x_W, x_N) \notin A$.

C. Linear combinations via discrete derivatives

In the following, we attain the two goals defined above by constructing a gate $\Gamma_\uparrow$ from Γ and considering the difference $\mathrm{Sig}(\Gamma_\uparrow) - \mathrm{Sig}(\Gamma)$. The gate $\Gamma_\uparrow$ is obtained from Γ by adding dummy rows of vertices with signature PRE, and this allows us to obtain $\mathrm{Sig}(\Gamma_\uparrow)$ by a simple substitution on the indeterminates of $\mathrm{Sig}(\Gamma)$.

Definition VI.3. We define a *dummy gate* as in Figure 4: Starting from a vertex with signature PRE, add several vertices of signature $\mathtt{HW}_{=1}$ to its western and eastern dangling edges to force these edges to be inactive, as shown in the left part of the figure. We then define a *dummy row* by arranging n dummy gates horizontally as shown in the right part of the figure.

Starting from Γ, define a gate $\Gamma_\uparrow$ by adding a dummy row above the row $(1, \star)$, and a dummy row below the row $(n, \star)$, as shown in Figure 4. We connect apex a_1 to the dangling edge 5 of each dummy gate, and a_2 to the dangling edge 6.

Furthermore, we define algebraic manipulations on multivariate polynomials that correspond to adding dummy rows as described above.

Definition VI.4. Let p be any multivariate polynomial over the indeterminates $\alpha_{u,v}$ and $\beta_{u,v}$ for $u, v \in [n]$. Write $x \leftarrow y$ for the operation of substituting x with y in p. Then we define $p_\uparrow$ to be the polynomial obtained from p after performing the substitutions $\alpha_{u,v} \leftarrow \alpha_{u,v} + 1$ and $\beta_{u,v} \leftarrow \beta_{u,v} + 1$ for all $u, v \in [n]$. We also define the following "discrete derivative" operator D on such polynomials p by declaring $D(p) := p_\uparrow - p$.

Lemma VI.5. *We have* $\mathrm{Sig}(\Gamma_\uparrow) = (\mathrm{Sig}(\Gamma))_\uparrow$*, and in particular, we have* $D(\mathrm{Sig}(\Gamma)) = \mathrm{Sig}(\Gamma_\uparrow) - \mathrm{Sig}(\Gamma)$.

Note that $D(p)$ indeed resembles a derivative: We have linearity by $D(p+q) = D(p)+D(q)$, and applying D to a polynomial p of degree d gives one of degree $d-1$. We will use D to effect two useful modifications on the polynomials in (11)-(13), and thus ultimately on $\mathrm{Sig}(\Gamma)$. These correspond to the two goals described at the end of Section VI-B.

1) Concerning the first goal, our choice of D ensures that "unwanted" polynomials vanish under D. For instance, for all $u, v, w \in [n]$, the polynomial $p_{u,v,w}$ from (12) maps to $D(p_{u,v,w}) = 0$. By our calculation of $\mathrm{Sig}(\Gamma)$ in Lemma VI.2, this implies that $D(\mathrm{Sig}(\Gamma))$ vanishes on assignments x with $x_N \neq x_S$ and $(x_W, x_N) \notin A$ and $(x_W, x_S) \notin A$. The other unwanted cases will be handled by similar arguments.
2) Under the operator D, linear terms, such as $\alpha_{u,v}$ for $u, v \in [n]$, are mapped to $D(\alpha_{u,v}) = (\alpha_{u,v} + 1) - \alpha_{u,v} = 1$. This helps us to attain the second goal, since the original terms depend on the concrete values of $\alpha_{u,v}$ in A, whereas the constants resulting from a derivative do not. It will also turn out that only linear terms are relevant.

In the following, we show that $D(\mathrm{Sig}(\Gamma))$ essentially realizes the function g_κ, up to some additive term on assignments x with φ_{prop}. This allows us to write g_κ as a linear combination of the matchgate signatures $\mathrm{Sig}(\Gamma_\uparrow)$ and $\mathrm{Sig}(\Gamma)$. As a technical requirement, we use Lemma II.7 to ensure that the set A in the definition of $\Gamma = \Gamma(A)$ is horizontally balanced.

Lemma VI.6. *Using Lemma II.7, assume the existence of a number $T \in \mathbb{N}$ such that A features exactly T elements of type $(u, \star)$, for all $u \in [n]$. Let $\Gamma = \Gamma(A)$ and write $D := D(\mathrm{Sig}(\Gamma)) = \mathrm{Sig}(\Gamma_\uparrow) - \mathrm{Sig}(\Gamma)$. For $x \in \{0,1\}^{4n}$ satisfying φ_{one}, we then have*

$$D(x) = \begin{cases} 0 & \text{if } \neg\varphi_{prop}(x) \\ \begin{cases} n - 2T - 2 & (x_W, x_N) \notin A \\ n - 2T + 2 & (x_W, x_N) \in A \end{cases} & \text{if } \varphi_{prop}(x) \end{cases}$$

Write $S := n - 2T - 2$. Then the following linear combination realizes the signature g_κ:

$$g_\kappa = \frac{D - S \cdot \mathrm{Sig}(\Phi)}{4} = \frac{\mathrm{Sig}(\Gamma_\uparrow) - \mathrm{Sig}(\Gamma) - S \cdot \mathrm{Sig}(\Phi)}{4}.$$

Note that the constituent gates $\Gamma_\uparrow$, Γ and Φ all have at most two apices and feature only edge-weights from the set $\{-1, 1\}$.

We can finally complete the proof of Theorem I.3. Recall that we reduce from $\oplus$GridTiling.

Proof of Theorem I.3. Let $\mathcal{A} = (n, k, \mathcal{C}, \mathcal{T})$ be an instance for $\oplus$GridTiling. For the lower bound under $\oplus$ETH, we may assume $|\mathcal{C}| = \mathcal{O}(k)$ by Theorem II.6. Furthermore, by Lemma II.7, we may assume to be given a number $T \in \mathbb{N}$ such that $|\mathcal{T}(\kappa) \cap (u, \star)| = T$ for all $\kappa \in \mathcal{C}$ and $u \in [n]$. Recall Definition V.2 and Lemma V.3 of Section V-A: These allow us to compute a signature graph G with signatures f_κ at $\kappa \in [k]^2 \setminus \mathcal{C}$ and signatures g_κ at $\kappa \in \mathcal{C}$ such that $\#\mathrm{GridTiling}(\mathcal{A}) = \mathrm{Holant}(G)$.

As shown in Lemma V.5, we can realize f_κ by the planar matchgate Φ on edge-weights $\{-1, 1\}$. Furthermore, as shown in Lemma VI.6, we can realize g_κ for each $\kappa \in \mathcal{C}$ as the linear combination of three 2-apex matchgates on edge-weights $\{-1, 1\}$: Let $\Gamma_\kappa := \Gamma(\mathcal{T}(\kappa))$ be as in Definition VI.1, and let $\Gamma_{\kappa,\uparrow}$ be obtained from Γ_κ as in Definition VI.3. Then, similarly to the proof of Theorem I.2, we obtain with Lemma VI.6 and Lemma III.6 about the linear combinations of signatures that

$$4^{|\mathcal{C}|} \cdot \mathrm{Holant}(G) = \sum_{\omega : \mathcal{C} \to [3]} (-1)^{d(\omega)} \cdot (-S)^{e(\omega)} \cdot \mathrm{perm}(H_\omega). \tag{14}$$

Here, for each $\omega : \mathcal{C} \to [3]$, the number $d(\omega)$ is defined to be the number of 2-entries in ω, and $e(\omega)$ is the number of 3-entries. The graph H_ω is obtained as follows: For $\kappa \in [k]^2 \setminus \mathcal{C}$, insert the matchgate Φ at the cell vertex c_κ. For all $\kappa \in \mathcal{C}$, insert $\Gamma_{\kappa,\uparrow}$ or Γ_κ or Φ at c_κ if $\omega(\kappa)$ is 1 or 2 or 3, respectively.

Define $M := 2^{2|\mathcal{C}|}$. Given an oracle for computing $\mathrm{perm}(H_\omega)$ modulo $2M$ for all ω, we can compute the right-hand side of (14) using arithmetic in $\mathbb{Z}/2M\mathbb{Z}$. We then obtain

$$M \cdot \mathrm{Holant}(G) = M \cdot \#\mathrm{GridTiling}(\mathcal{A}) \equiv_{2M} \begin{cases} M & \text{if } \#\mathrm{GridTiling}(\mathcal{A}) \text{ odd}, \\ 0 & \text{if } \#\mathrm{GridTiling}(\mathcal{A}) \text{ even}. \end{cases}$$

Each graph H_ω is bipartite, has at most $2|\mathcal{C}| = \mathcal{O}(k)$ apices, and the computation is modulo $2M = 2^{\mathcal{O}(k)}$. We have thus shown a parameterized Turing reduction from $\oplus$GridTiling to the evaluation of the permanent on $\mathcal{O}(k)$-apex graphs modulo $2^{\mathcal{O}(k)}$. Together with Theorem II.6, the theorem follows. □

ACKNOWLEDGMENT

Radu Curticapean is supported by the ERC Starting Grant PARAMTIGHT, No. 280152. Mingji Xia is supported by China National 973 program 2014CB340301, China Basic Research Program (973) Grant 2014CB340302, NSFC 61003030 and NSFC 61170073. We wish to thank an anonymous reviewer for extensive and helpful comments on the submitted version.

References

[1] L. G. Valiant, "The complexity of computing the permanent," *Theoret. Comput. Sci.*, vol. 8, no. 2, pp. 189–201, 1979.

[2] ——, "The complexity of enumeration and reliability problems," *SIAM J. Comput.*, vol. 8, no. 3, pp. 410–421, 1979. [Online]. Available: http://dx.doi.org/10.1137/0208032

[3] H. N. V. Temperley and M. E. Fisher, "Dimer problem in statistical mechanics - an exact result," *Philosophical Magazine*, vol. 6, no. 68, pp. 1478–6435, 1961.

[4] P. W. Kasteleyn, "The statistics of dimers on a lattice: I. The number of dimer arrangements on a quadratic lattice," *Physica*, vol. 27, no. 12, pp. 1209 – 1225, 1961. [Online]. Available: http://www.sciencedirect.com/science/article/B6X42-46CBXYS-H/2/e0d1c417527cca4e7e20eeb185ab8586

[5] ——, "Graph Theory and Crystal Physics," in *Graph Theory and Theoretical Physics*. Academic Press, 1967, pp. 43–110.

[6] M. Agrawal, "Determinant versus permanent," in *Proceedings of the 25th International Congress of Mathematicians, ICM 2006*, vol. 3, 2006, pp. 985–997.

[7] R. Raz, "Multi-linear formulas for permanent and determinant are of super-polynomial size," *J. ACM*, vol. 56, no. 2, 2009. [Online]. Available: http://doi.acm.org/10.1145/1502793.1502797

[8] J.-Y. Cai, X. Chen, and D. Li, "Quadratic lower bound for permanent vs. determinant in any characteristic," *Computational Complexity*, vol. 19, no. 1, pp. 37–56, 2010. [Online]. Available: http://springerlink.metapress.com/content/8314445277 1m3k01/

[9] L. G. Valiant, "Completeness classes in algebra," in *Proceedings of the eleventh annual ACM symposium on Theory of computing*, ser. STOC '79. New York, NY, USA: ACM, 1979, pp. 249–261. [Online]. Available: http://doi.acm.org/10.1145/800135.804419

[10] M. Jerrum, A. Sinclair, and E. Vigoda, "A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries," *J. ACM*, vol. 51, no. 4, pp. 671–697, 2004. [Online]. Available: http://doi.acm.org/10.1145/1008731.1008738

[11] M. Jerrum and A. Sinclair, "Polynomial-time approximation algorithms for the ising model," *SIAM J. Comput.*, vol. 22, no. 5, pp. 1087–1116, 1993. [Online]. Available: http://dx.doi.org/10.1137/0222066

[12] A. Galluccio and M. Loebl, "On the theory of Pfaffian orientations. I. Perfect matchings and permanents." *Electronic Journal of Combinatorics*, vol. 6, 1999.

[13] G. Tesler, "Matchings in graphs on non-orientable surfaces," *J. Comb. Theory, Ser. B*, vol. 78, no. 2, pp. 198–231, 2000. [Online]. Available: http://dx.doi.org/10.1006/jctb.1999.1941

[14] T. Regge and R. Zecchina, "Combinatorial and topological approach to the 3d ising model," *Journal of Physics A: Mathematical and General*, vol. 33, no. 4, pp. 741–761, 2000.

[15] R. G. Downey and M. R. Fellows, "Parameterized computational feasibility," in *Feasible Mathematics II*. Birkhäuser Boston.

[16] ——, *Parameterized Complexity*, ser. Monographs in Computer Science. Springer, 1999. [Online]. Available: http://dx.doi.org/10.1007/978-1-4612-0515-9

[17] J. Flum and M. Grohe, "The parameterized complexity of counting problems," *SIAM Journal on Computing*, no. 4, pp. 892–922, 2004.

[18] C. McCartin, "Parameterized counting problems," *Ann. Pure Appl. Logic*, vol. 138, no. 1-3, pp. 147–182, 2006. [Online]. Available: http://dx.doi.org/10.1016/j.apal.2005.06.010

[19] J. Flum and M. Grohe, *Parameterized Complexity Theory*. Springer, 2006.

[20] R. Curticapean, "Counting matchings of size k is #W[1]-hard," in *ICALP (1)*, 2013, pp. 352–363.

[21] R. Curticapean and D. Marx, "Complexity of counting subgraphs: Only the boundedness of the vertex-cover number counts," in *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, 2014, pp. 130–139. [Online]. Available: http://dx.doi.org/10.1109/FOCS.2014.22

[22] P. Dagum and M. Luby, "Approximating the permanent of graphs with large factors," *Theor. Comput. Sci.*, vol. 102, no. 2, pp. 283–305, 1992. [Online]. Available: http://dx.doi.org/10.1016/0304-3975(92)90234-7

[23] C. Little, "An extension of Kasteleyn's method of enumerating the 1-factors of planar graphs," in *Combinatorial Mathematics*, ser. LNCS, 1974, pp. 63–72.

[24] V. V. Vazirani, "NC algorithms for computing the number of perfect matchings in $K_{3,3}$-free graphs and related problems," *Inf. Comput.*, vol. 80, no. 2, pp. 152–164, 1989.

[25] S. Straub, T. Thierauf, and F. Wagner, "Counting the number of perfect matchings in K_5-free graphs," in *IEEE 29th Conference on Computational Complexity, CCC 2014, Vancouver, BC, Canada, June 11-13, 2014*, 2014, pp. 66–77. [Online]. Available: http://dx.doi.org/10.1109/CCC.2014.15

[26] R. Curticapean, "Counting perfect matchings in graphs that exclude a single-crossing minor," *CoRR*, vol. abs/1406.4056, 2014. [Online]. Available: http://arxiv.org/abs/1406.4056

[27] N. Robertson and P. D. Seymour, "Graph minors. XVI. Excluding a non-planar graph," *Journal of Combinatorial Theory, Series B*, vol. 89, no. 1, pp. 43 – 76, 2003. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S009589560300042X

[28] ——, "Graph minors. XX. Wagner's conjecture," *J. Comb. Theory, Ser. B*, vol. 92, no. 2, pp. 325–357, 2004. [Online]. Available: http://dx.doi.org/10.1016/j.jctb.2004.08.001

[29] D. Marx and M. Pilipczuk, "Everything you always wanted to know about the parameterized complexity of subgraph isomorphism (but were afraid to ask)," in *31st International Symposium on Theoretical Aspects of Computer Science (STACS 2014), STACS 2014, March 5-8, 2014, Lyon, France*, 2014, pp. 542–553. [Online]. Available: http://dx.doi.org/10.4230/LIPIcs.STACS.2014.542

[30] E. D. Demaine, M. Hajiaghayi, and K. Kawarabayashi, "Approximation algorithms via structural results for apex-minor-free graphs," in *Automata, Languages and Programming, 36th International Colloquium, ICALP 2009, Rhodes, Greece, July 5-12, 2009, Proceedings, Part I*, 2009, pp. 316–327.

[31] E. D. Demaine, M. T. Hajiaghayi, and K. Kawarabayashi, "Algorithmic graph minor theory: Decomposition, approximation, and coloring," in *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), 23-25 October 2005, Pittsburgh, PA, USA, Proceedings*, 2005, pp. 637–646. [Online]. Available: http://dx.doi.org/10.1109/SFCS.2005.14

[32] E. Demaine, M. Hajiaghayi, and D. Thilikos, "Exponential speedup of fixed-parameter algorithms for classes of graphs excluding single-crossing graphs as minors," *Algorithmica*, vol. 41, no. 4, pp. 245–267, 2005.

[33] ——, "1.5-approximation for treewidth of graphs excluding a graph with one crossing as a minor," in *APPROX*, 2002, pp. 67–80.

[34] F. V. Fomin, E. D. Demaine, and M. T. Hajiaghayi, "Bidimensionality," in *Encyclopedia of Algorithms*, 2015.

[35] D. Marx and I. Schlotter, "Obtaining a planar graph by vertex deletion," *Algorithmica*, vol. 62, no. 3-4, pp. 807–822, 2012. [Online]. Available: http://dx.doi.org/10.1007/s00453-010-9484-z

[36] A. Björklund, T. Husfeldt, and I. Lyckberg, "Computing the permanent modulo a prime power," unpublished, and personal communication.

[37] A. Björklund, H. Dell, and T. Husfeldt, "The parity of set systems under random restrictions with applications to exponential time problems," in *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, 2015, pp. 231–242.

[38] D. Marx, "Can you beat treewidth?" *Theory of Computing*, vol. 6, no. 1, pp. 85–112, 2010.

[39] J.-Y. Cai and P. Lu, "Holographic algorithms: From art to science," in *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, ser. STOC '07. New York, NY, USA: ACM, 2007, pp. 401–410. [Online]. Available: http://doi.acm.org/10.1145/1250790.1250850

[40] J.-Y. Cai and V. Choudhary, "Valiant's holant theorem and matchgate tensors," in *In Proceedings of TAMC 2006: Lecture Notes in Computer Science*, 2006, pp. 248–261.
[41] J.-Y. Cai and A. Gorenstein, "Matchgates revisited," *CoRR*, vol. abs/1303.6729, 2013.
[42] H. Guo, P. Lu, and L. G. Valiant, "The complexity of symmetric Boolean parity Holant problems," *SIAM Journal on Computing*, vol. 42, no. 1, pp. 324–356, 2013. [Online]. Available: http://dx.doi.org/10.1137/100815530
[43] J.-Y. Cai, H. Guo, and T. Williams, "A complete dichotomy rises from the capture of vanishing signatures: extended abstract," in *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, 2013, pp. 635–644. [Online]. Available: http://doi.acm.org/10.1145/2488608.2488687
[44] R. Diestel, *Graph Theory, 4th Edition*, ser. Graduate texts in mathematics. Springer, 2012, vol. 173.
[45] R. Impagliazzo and R. Paturi, "On the complexity of k-SAT," *Journal of Computer and System Sciences*, vol. 62, no. 2, pp. 367–375, 2001.
[46] R. Impagliazzo, R. Paturi, and F. Zane, "Which problems have strongly exponential complexity?" *J. Computer and Sys. Sci.*, vol. 63, no. 4, pp. 512–530, 2001.
[47] H. Dell, T. Husfeldt, D. Marx, N. Taslaman, and M. Wahlen, "Exponential time complexity of the permanent and the Tutte polynomial," *ACM Transactions on Algorithms*, vol. 10, no. 4, p. 21, 2014. [Online]. Available: http://doi.acm.org/10.1145/2635812
[48] J. Chen, B. Chor, M. Fellows, X. Huang, D. W. Juedes, I. A. Kanj, and G. Xia, "Tight lower bounds for certain parameterized np-hard problems," *Inf. Comput.*, vol. 201, no. 2, pp. 216–231, 2005. [Online]. Available: http://dx.doi.org/10.1016/j.ic.2005.05.001
[49] J. Chen, X. Huang, I. A. Kanj, and G. Xia, "Strong computational lower bounds via parameterized complexity," *Journal of Computer and System Sciences*, vol. 72, no. 8, pp. 1346 – 1367, 2006. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0022000006000675
[50] L. G. Valiant and V. V. Vazirani, "NP is as easy as detecting unique solutions," *Theoretical Computer Science*, vol. 47, pp. 85–93, 1986.
[51] C. Calabro, R. Impagliazzo, V. Kabanets, and R. Paturi, "The complexity of unique k-SAT: An isolation lemma for k-CNFs," *J. Comput. Syst. Sci.*, vol. 74, no. 3, pp. 386–393, 2008. [Online]. Available: http://dx.doi.org/10.1016/j.jcss.2007.06.015
[52] D. Marx, "A tight lower bound for planar multiway cut with fixed number of terminals," in *Automata, Languages, and Programming - 39th International Colloquium, ICALP 2012, Warwick, UK, July 9-13, 2012, Proceedings, Part I*, 2012, pp. 677–688.
[53] L. G. Valiant, "Holographic algorithms," *SIAM J. Comput.*, vol. 37, no. 5, pp. 1565–1594, 2008.
[54] J.-Y. Cai, P. Lu, and M. Xia, "Holographic algorithms by Fibonacci gates and holographic reductions for hardness," in *Proceedings of the 2008 49th Annual IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society, 2008, pp. 644–653. [Online]. Available: http://portal.acm.org/citation.cfm?id=1470582.1470613