The Precise Complexity of Finding Rainbow Even Matchings*

Martin Loebl[0000−0001−7968−0376]

Dept. of Applied Mathematics, Charles University, Prague, Czech Republic.
loebl@kam.mff.cuni.cz

Abstract. A progress in complexity lower bounds might be achieved by studying problems where a very precise complexity is conjectured. In this note we propose one such problem: Given a planar graph on n vertices and disjoint pairs of its edges $p_1, \dots, p_g$, perfect matching M is Rainbow Even Matching (REM) if $|M \cap p_i|$ is even for each $i = 1, \dots, g$. A straightforward algorithm finds a REM or asserts that no REM exists in $2^g \times \text{poly}(n)$ steps and we conjecture that no deterministic or randomised algorithm has complexity asymptotically smaller than 2^g. Our motivation is also to pinpoint the curse of dimensionality of the Max-Cut problem for graphs embedded into orientable surfaces: a basic problem of statistical physics.

Keywords: matching · max cut · exponential time hypothesis · Ising partition function.

1 Introduction

Given a graph $G = (V, E)$, a set of edges $M \subseteq E$ is called *perfect matching* if the graph (V, M) has degree one at each vertex. In this paper we introduce and study the following matching problems which, as far as we know, were not studied before.

Given a graph $G = (V, E)$ and disjoint pairs of its edges $p_1, \dots, p_g$, we say that a perfect matching M is a Rainbow Even Matching (REM) if $|M \cap p_i|$ is even for each $i = 1, \dots, g$. For example, let C be a cycle of length 8 consisting of consecutive edges $e_1, e_2, \dots, e_8$. If $g \geq 1$ and $p_1 = \{e_1, e_2\}$ then there is no REM and if $g = 3$ and $p_1 = \{e_1, e_3\}, p_2 = \{e_2, e_4\}, p_3 = \{e_5, e_6\}$ then both perfect matchings of C are REM. We consider the following problems:

1. Decision Rainbow Even Matching problem (DREM): Given a planar graph G on n vertices and disjoint pairs of edges $p_1, \dots, p_g$, decide if there is a REM.

2. Enumeration Rainbow Even Matching problem (EREM): Given a planar graph G on n vertices and disjoint pairs of edges $p_1, \dots, p_g$, calculate the number of REMs.

* The author was partially supported by the H2020-MSCA-RISE project CoSP- GA No. 823748.

3. If an integer weight function is given on the edge-set of the graph G then DREM has a natural weighted version, denoted by **OptDREM**, to find the maximum total weight of a REM, and EREM is turned into the problem denoted by **GenREM** to find the generating function of weighted REMs.

There is a straightforward algorithm of complexity $2^g\text{poly}(n)$ to solve OptDREM: For each $S \subset \{1, \dots, g\}$ we find a maximum weight extension of the the set $\cup_{i \in S} p_i$ into a perfect matching by edges of $E \setminus \cup_{i \leq g} p_i$. The weighted perfect matching algorithm does it.

There is also a straightforward algorithm of complexity $2^g\text{poly}(n)$ to solve GenREM: For each $S \subset \{1, \dots, g\}$ we calculate the generating function of the REMs which contain all edges of $\cup_{i \in S} p_i$ and no edge of $\cup_{i \notin S} p_i$. This can be done by the *method of Kasteleyn orientations* briefly introduced in subsection 1.3.

Main contribution.

- We propose that the above standard algorithms are in fact optimal. Our Frustration Conjecture 1 below states that up to a polynomial factor the *precise complexity* of OptDREM with edge-weight in $\{-1, 0, 1\}$ is 2^g. This is more tight complexity specification than the *Strong Exponential Time Hypothesis*.
- We show that refuting the Frustration Conjecture 1 implies that in the class of graphs where the crossing number is equal to the genus, the complexity of the Max-Cut problem is smaller than the *additive determinantal complexity* of cuts enumeration. At present, no natural class of embedded graphs with this property is known.

1.1 The Exponential Time Hypothesis

The *Exponential Time Hypothesis (ETH)* is an unproven computational hardness assumption that was formulated by Impagliazzo and Paturi [7]. For each k let s_k be the infimum of reals s for which there exists an algorithm solving $k-$SAT in time $O(2^{sn})$, where n is the number of variables. ETH states that for each $k > 2$, $s_k > 0$. We note that $2-$SAT can be solved in polynomial time. In the same paper [7], the authors prove the *Sparsification Lemma* which implies that ETH is equivalent to a potential strengthening of ETH where the $k-$SAT instances have the number of clauses bounded from above by $c_k n$ for some constant c_k; n denotes the number of the variables.

The ETH was strengthened by Impagliazzo, Paturi and Zane [8] to the *Strong Exponential Time Hypothesis (SETH)*: For all $d < 1$ there is a k such that k-SAT cannot be solved in $O(2^{dn})$ time. No Sparsification Lemma is known for SETH.

Both ETH and SETH have a very natural role: they are used to argue that known algorithms are probably optimal.

I believe that the method of Kasteleyn orientations provides optimal algorithms for the Max-Cut problem in the classes of embedded graphs.

My motivation for introducing REM has been to pinpoint this 'curse of dimensionality' by a problem formulated with no reference to the geometry.

Conjecture 1 (Frustration Conjecture). No deterministic or randomised algorithm can solve OptDREM with the edge-weights from $\{-1, 0, 1\}$ in asymptotically less than 2^g steps.

1.2 Justification for the Frustration Conjecture

An exponential lower bound for DREM is simply implied by ETH, see Corollary 1. Next, Theorem 2 connects the Frustration Conjecture 1 to the additive determinantal complexity of cuts enumeration.

A well-established way to approach matching problems is to determine whether some specific coefficient of the generating function of the perfect matchings (with suitable substitutions) is non-zero. This can be achieved because of the *Isolation Lemma*, see [13], by calculating a single Pfaffian of a matrix where the entries are monomials in possibly more than one variable. The Pfaffian is a determinant type expression which can be computed with essentially the same complexity as that of the determinant (of the same matrix). The complexity of calculating the determinant of matrices with polynomial entries essentially depends on the number of the variables.

After many failed attempts to use this machinery to disprove the Frustration Conjecture I am convinced that this approach will not beat the 2^g lower bound. However, I do not have at present a general theorem of this nature, only some partial results.

We can reduce, in a simple way suggested by Bruno Loff, (1 in 3)-SAT to DREM showing DREM is NP-complete.

Theorem 1. *DREM is an NP-complete problem.*

Proof. The reduction of (1 in 3)-SAT to DREM is best explained by an example. If the input of (1 in 3)-SAT is $(x_1 \vee \neg x_2 \vee x_3) \wedge (x_2 \vee x_1 \vee x_4)$ where the first clause is denoted by C_1 and the second clause by C_2 then the input graph for the corresponding DREM is depicted in Figure 1, with $g = 2$ and $p_1^1 = \{e_1^1, e_2^1\}, p_1^2 = \{e_1^2, e_2^2\}$. This simply generalises.

Let $x_1, \ldots, x_n$ be the variables and let $C_1, \ldots, C_m$ be the clauses of a (1 in 3)-SAT input. (1) With each clause C_j we associate a copy $S(j)$ of the star with three leaves. (2) Let $x(i, j)$ denote the appearance of variable x_i in clause C_j. (3) If $x(i, j)$ is equal to x_i then let $P(i, j)$ be a copy of the path of three edges. (4) If $x(i, j)$ is equal to $\neg x_i$ then let $P(i, j)$ be a copy of the path of five edges. (5) Let variables $x_{i_1}, x_{i_2}, x_{i_3}$ appear in clause C_j. Then we identify the two leaves of $P(i_1, j)$ $(P(i_2, j), P(i_3, j)$ respectively) with two leaves of $S(j)$ as indicated in Figure 1. (6) Finally we specify $g = 3m - n$ disjoint pairs of edges: Let $i \leq n$ and let $x(i, j_1), \ldots, x(i, j_k)$ be all the appearances of veriable x_i. For $l = 1, \ldots, k$ let $e(i, j_l)$ be an edge adjacent to the middle edge of $P(i, j_l)$. For each $i \leq n$ and $l \in \{1, \ldots, k-1\}$ we will have edge-pair $p_l^i = \{e(i, j_l), e(i, j_{l+1})\}$.

(7) These pairs assure the following: Let M be a REM and $1 \leq i \leq n$. Then the middle edge of each $P(i, j)$ belongs to M or the middle edge of no

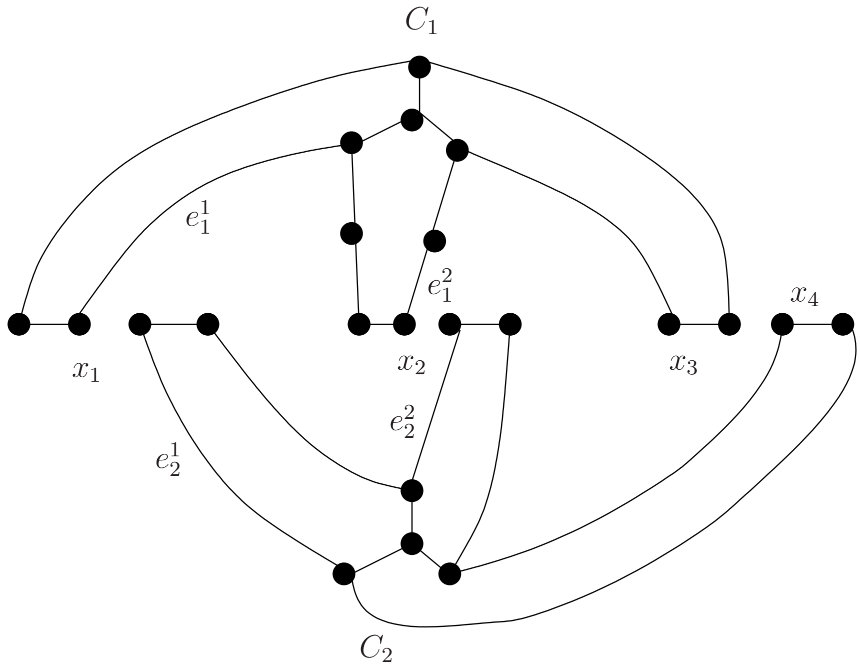

Fig. 1.

$P(i,j)$ belongs to M. This simply implies that there is a REM iff there is a (1 in 3)-satisfying assignement.

Corollary 1. *Let D be the infimum of reals d for which there exists an algorithm solving DREM in time $O(2^{dg})$, where g is the number of the input pairs of edges. Let us assume that the Exponential time hypothesis holds. Then $D > 0$.*

Proof. We first note that 3-SAT with n variables and m clauses can be reduced to (1 in 3)-SAT with $n + 6m$ variables and $5m$ clauses by a construction of Schaefer [14]. By the discussion in 1.1 we can assume that $m \leq c_3 n$. After this reduction we use the construction of the proof of Theorem 1.

1.3 Kasteleyn Orientations and Optimisation by Enumeration

Let me state a curious phenomenon: There is a strongly polynomial algorithm to solve the Max-Cut problem in the planar graphs based on a reduction to the weighted perfect matching problem, see e.g. [10].

For the graphs of fixed genus $g \geq 1$ the situation is different: There is a weakly polynomial algorithm by Galluccio and Loebl ([4]; see also [5], [6]); it was implemented several times and applied in extensive statistical physics calculations (see [12]). Recently other related algorithms based on the Valiant's theory [16] of holographic algorithms appeared (see [1], [3]). All presently known approaches are of enumeration nature even for the class of the toroidal square grids. The weakly polynomial **optimisation by enumeration method** of [4] is as follows:

1. Let $G = (V, E)$ be a graph. A set of edges $E' \subseteq E$ is called *even* if each degree of the graph (V, E') is even. A set of edges $C \subseteq E$ is called an *edge-cut* of G, if there is a $V' \subseteq V$ so that $C = \{e \in E : |e \cap V'| = 1\}$. The MAX-CUT problem, one of the basic optimisation problems, asks for the maximum size of an edge-cut in the input graph G, or, if weights on the edges are given, for the maximum total weight of an edge-cut.

2. If a weight-function $w : E \to \mathbf{R}$ and a set S of subsets of E are given then the *generating function of* S is defined as

$$\mathcal{F}(G, w, x) = \sum_{A \in S} \prod_{e \in A} x^{w(e)}.$$

3. The generating function of the edge-cuts is simply equivalent to the Ising partition function of the same graph, and it can be computed from the generating function of the even sets by a theorem of Van der Waender (for the definitions, theorems and their proofs see e.g. [10]).

4. The generating function of the even sets can be computed by the Fisher construction described in 2.3 as the generating function of the perfect matchings of a modified graph.

5. The seminal technical proposition was formulated by Kasteleyn [9] and proved by Galluccio, Loebl [4] and independently by Tesler [15]:

The generating function of perfect matchings of a graph of genus g can be efficiently written as a linear combination of 2^{2g} Pfaffians. Pfaffians are determinant type expressions that can be computed efficiently by a variant of the Gaussian elimination. Cimasoni and Reshetikhin [2] provided a beautiful interpretation of the formula which then became known as the Arf invariant formula.

6. Summarising, the weakly polynomial algorithm solving the MAX-CUT problem for the graphs of genus g by Galluccio and Loebl consists in calculating 2^{2g} Pfaffians and produces the complete generating function of the edge-cuts of the embedded graph.

1.4 Additive determinantal complexity

A recent result of Loebl and Masbaum [11] indicates that this might be optimum for the cuts enumeration. It is shown by Loebl and Masbaum in [11] that, if we want to enumerate the edge-cuts of each possible size of an input graph G of genus g, then in a strongly restricted setting called *additive determinantal complexity* the number of the Pfaffian calculations cannot be smaller than 2^{2g}.

This leads to a question: Is there an algorithm for solving the MAX-CUT problem in (a natural subclass of) the embedded graphs, whose complexity *beats* the additive determinantal complexity of the cuts enumeration? At present no such algorithm for a natural subclass of embedded graphs is known.

I believe that the answer to this question is NO and in Theorem 2 below we present a partial result. We show that the Frustration Conjecture implies that for the class of embedded graphs where the crossing number is equal to the genus, there is no algorithm to solve the MAX-CUT problem whose complexity beats the additive determinantal complexity bound. The proof of Theorem 2 is included in Section 2.

Theorem 2. *Let G be a graph with n vertices and embedded to the plane with g crossings. One can efficiently construct planar graph G' with edge-weights in $\{-1,0,1\}$ and a set of $2g$ disjoint pairs of edges of G' so that finding the maximum size of an edge-cut in G is polynomial time reducible to determining the maximum weight of a REM in G'.*

Acknowledgement. This project initially started as a joint work with Marcos Kiwi. I would like to thank Marcos for many helpful discussions.

2 Edge-Cuts in Embedded Graphs

Let $G = (V, E)$ be a graph. A set of edges $E' \subseteq E$ is called *even* if each degree of the graph (V, E') is even. A set of edges $C \subseteq E$ is called an *edge-cut* of G, if there is a $V' \subseteq V$ so that $C = \{e \in E : |e \cap V'| = 1\}$. The MAX-CUT problem, one of the basic optimisation problems, asks for the maximum size of an edge-cut in the input graph G, or, if weights on the edges are given, for the maximum total weight of an edge-cut.

2.1 Surfaces

We recall the following standard description of a genus g surface S_g with one boundary component (we follow [10], [11]). (We reserve the notation Σ_g for a closed surface of genus g.)

Definition 1. *A* 1-highway *(see Figure 2) is a surface $\bar{S}_g$ which consists of a base polygon R_0 and bridges $R_1, \ldots, R_{2g}$, where*

- *R_0 is a convex $4g$-gon with vertices $a_1, \ldots, a_{4g}$ numbered clockwise.*
- *Each R_{2i-1} is a rectangle with vertices $x(i,1), \ldots, x(i,4)$ numbered clockwise and glued to R_0. Edges $[x(i,1), x(i,2)]$ and $[x(i,3), x(i,4)]$ of R_{2i-1} are identified with edges $[a_{4(i-1)+1}, a_{4(i-1)+2}]$ and $[a_{4(i-1)+3}, a_{4(i-1)+4}]$ of R_0, respectively.*
- *Each R_{2i} is a rectangle with vertices $y(i,1), \ldots, y(i,4)$ numbered clockwise and glued to R_0. Edges $[y(i,1), y(i,2)]$ and $[y(i,3), y(i,4)]$ of R_{2i-1} are identified with edges $[a_{4(i-1)+2}, a_{4(i-1)+3}]$ and $[a_{4(i-1)+4}, a_{4(i-1)+5}]$ of R_0, respectively. (Here, indices are considered modulo $4g$.)*

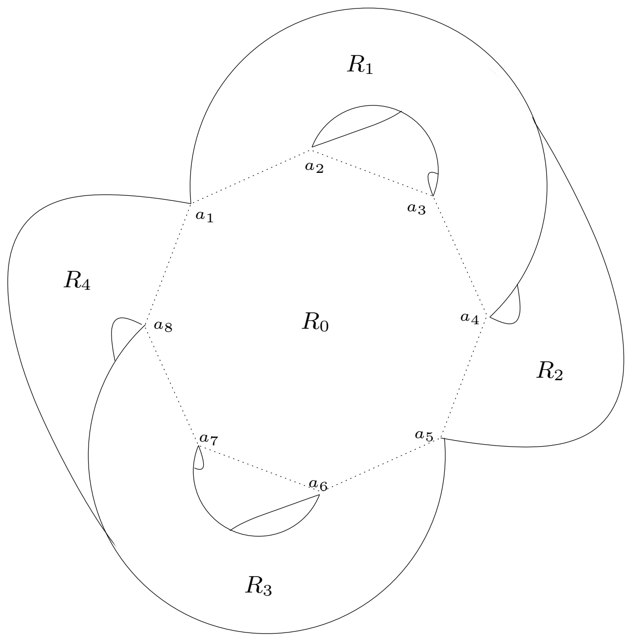

Fig. 2. A 1-highway.

Before proceeding, we point out a simple fact that we will soon exploit: the boundary of a 1-highway is isotopic to the boundary of a disk.

Now assume the graph G is embedded into a closed orientable surface Σ_g of genus g. We think of Σ_g as S_g union an additional disk δ glued to the boundary of S_g. By an isotopy of the embedding, we may assume that G does not meet the disk δ and that, moreover, all vertices of G lie in the interior of R_0.

We may also assume that the intersection of G with any of the rectangular bridges R_i consists of disjoint straight lines connecting the two sides of R_i which are glued to the base polygon R_0.

Next, follows the standard analogous description of a genus g surface S_g with more than one boundary component.

Definition 2. *A* highway surface S_g *is obtained from a 2-sphere Z with h disjoint polygons $R_0^1, \ldots, R_0^h$ specified, and h disjoint 1-highway surfaces $\bar{S}_{g_1}^1, \ldots, \bar{S}_{g_h}^h$, where $g = g_1 + \ldots + g_h$, by first identifying the base polygon of each $\bar{S}_{g_i}^i$ with the polygon R_0^i, and then by deletion of the interiors of these polygons R_0^i $(i = 1, \ldots, h)$.*

Now assume the graph G is embedded into a closed orientable surface Σ_g of genus g. We again think of Σ_g as S_g union h additional disks δ_i $(i = 1, \ldots, h)$, glued to the h boundaries of S_g. By an isotopy of the embedding, we may assume that G does not meet the disks δ_i's and that, moreover, no vertex of G lies in a bridge. We may also assume that the intersection of G with any of the rectangular bridges R_i^j consists of disjoint straight lines connecting the two sides of R_i^j which are glued to the base sphere Z.

2.2 Local non-planarity

We note that each embedding of a graph G into Σ_g defines its *geometric dual*, usually denoted by G^*, as follows: the vertices of G^* are the faces of the embedding of G and for each edge e of G there is an edge e^* of G^* connecting the faces which have e on their boundary. For example, each toroidal square grid is self-dual. We note that a dual can have loops and multiple edges.

We consider simultaneous embeddings of the graph and its geometric dual into Σ_g.

Definition 3. *Let $G = (V, E)$ be a graph. A* simultaneous embedding *of G into Σ_g consists of (1) an embedding N of graph G, and (2) an embedding N^* of the geometric dual $G^* = (V^*, E^*)$ of N. In addition, we require that (a) G is the geometric dual of N^*, (b) each vertex of G^* (of G respectively) is embedded in the face of N (N^* respectively) it represents, (c) each pair of dual edges e, e^* intersects exactly once, and N, N^* have no other intersections, and (d) both N, N^* are embeddings into $S_g \subseteq \Sigma_g$.*

For a collection of edges $S \subseteq E$ we denote by $S^* \subseteq E^*$ the collection of dual edges e^* such that $e \in S$.

Since a simultaneous embedding of G into Σ_g is by definition a subset of $S_g \subseteq \Sigma_g$, we will also call it *simultaneous embedding into S_g*.

We may also assume that the intersection of G with any of the rectangular bridges R_i^j consists of disjoint straight lines connecting the two sides of R_i^j which are glued to the base sphere Z.

Definition 4. *We recall that the intersection of an embedding of G in S_g with any of the rectangular bridges R_i^j of $S)g$ consists of disjoint straight lines connecting the two sides of R_i^j which are glued to the base sphere.*

Let G be embedded in S_g. An even set $E' \subset E$ of the edges of G of which crosses each bridge of S_g by an even number of disjoint straight lines will be called admissible*.*

A simultaneous embedding of G into S_g is called even *if it holds that $C \subseteq E$ is an edge-cut of G if and only if $C^* \subseteq E^*$ is an admissible even set of the embedding of G^*.*

A basic example of an even simultaneous embedding is a toroidal square grid and its geometric dual.

Definition 5. *We say that a simultaneous even embedding of a graph G into some S_g is* restricted *if E^* intersects each bridge by at most 2 disjoint straight lines.*

Definition 6. *We say that graph G belongs to class $\mathcal{C}_g$ if G is drawn to the plane with exactly g edge-crossings and for each crossing there is a planar disc where the drawing looks as depicted in Figure 3.*

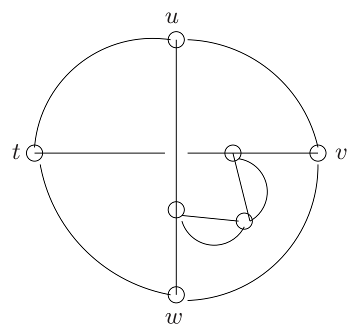

Fig. 3.

Theorem 3. *If $G \in \mathcal{C}_g$, then G admits a restricted even simultaneous embedding into S_g.*

Proof. We consider the simultaneous local embedding of the graph G as described in Figure 4. The embedding is clearly restricted. We need to show that the embedding is even.

We first observe that the set $\delta(v)$ of the edges of G incident with any vertex v of G satisfies that $\delta^*(v)$ intersects each bridge in an even number of segments. Since each edge-cut of G is the symmetric difference of some sets $\delta(v), v \in V$, we get: If C is an edge-cut of G, then C^* is admissible.

In order to prove that the embedding is even we need to show that each admissible set C^* of dual edges is a symmetric difference of faces of G^*; this implies that C is an edge-cut of G. We can assume that C^* has empty intersection with the bridges (depicted in Figure 4):

Consider the pair of bridges in Figure 4. There is a face F_1 of G^* with exactly 2 edges on the vertical bridge and no edge on the horizontal bridge, and also a face F_2 of G^* where the role of the two bridges is exchanged. We can use the symmetric difference of C^* with F_1 or F_2 to produce a new even set C_0^* which has empty intersection with each bridge. Moreover, if C_0^* is a symmetric difference of faces of G^* then so is C^*.

It follows that C^* is an even subset of an embedded planar subgraph of G^*. For the planar graphs, the boundaries of faces generate all even sets of edges by the symmetric difference operation. Hence the proof is finished if we show that each face F of this planar subgraph is a symmetric difference of the faces of G^*:

Indeed, such F is either a face of G^* itself, or it looks like the square of Figure 4 comprised of edges depicted as thick lines, which is the symmetric difference of the dual faces encircling the three unlabelled vertices of G of Figure 4.

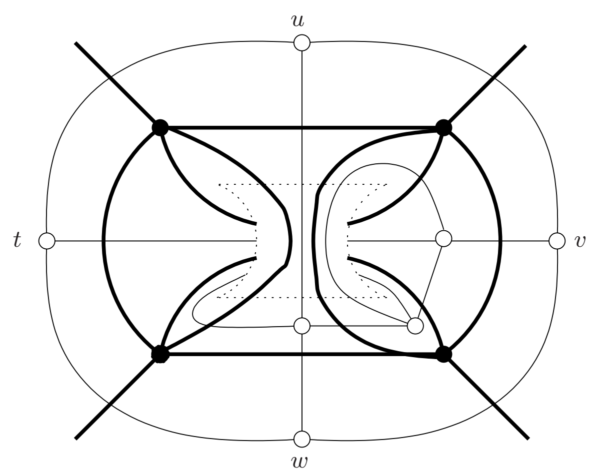

Fig. 4. Simultaneous embedding of graph $G \in \mathcal{C}_g$ near a crossing. There is one pair of bridges; the boundaries of the vertical bridge are depicted by dotted lines and the boundaries of the horizontal bridge are not depicted to simplify the presentation. The edges of G are depicted by normal lines and the dual edges are depicted by thick lines.

2.3 Proof of Theorem 2

We show that for a graph $G = (V, E)$ with n vertices and embedded to the plane with g edge crossings one can efficiently construct a planar graph $H = (W, E')$ with edge-weights in $\{-1, 0, 1\}$ and with $2g$ specified disjoint pairs of its edges so that the maximum size of an edge-cut in G is equal to the maximum weight of a REM in H. The construction goes as follows:

Step 1. We subdivide each edge of G near to each crossing; if $e \in E$ got subdivided into edges $e_1, \ldots, e_k$ which form the path $(e_1, \ldots, e_k)$ then we let the weight of e_1 equal to 1 and the weight of $e_2, \ldots, e_k$ equal to -1. The resulting weighted graph will be denoted by G_1. We note that the MAX-CUT problem in G is reduced to the weighted MAX-CUT problem in G_1.

Step 2. We add, for each edge crossing of G_1, the four edges of weight zero forming a $4-$cycle (denoted by $uvwt$ in Figure 4) and further one new vertex which we connect by four edges of weight zero to the two vertices near to this crossing added in Step 1 so that the resulting graph, which we denote by G_2, is in $\mathcal{C}_g$. We note that G_2 is uniquely determined and the weighted MAX-CUT problem in G_1 is reduced to the weighted MAX-CUT problem in G_2.

Step 3. We use Theorem 3. Let G_2^* be the dual from the restricted simultaneous even embedding of G_2 into S_g. The weight of each edge e^* of G_2^* is defined to be equal to the weight of the corresponding edge e of G_2. We specify $2g$ pairs $p_1, \ldots, p_{2g}$ of edges of G_2^*:

each pair consists of the two edges embedded on one of the $2g$ bridges of S_g (see Figure 4). We note that the weighted MAX-CUT problem for G_2 is reduced to the problem of finding maximum weight even set of G_2^* which contains an even number of elements of each pair $p_i, i = 1, \ldots, 2g$. Finally we note that G_2^* is planar.

Step 4: Fisher's construction. We transform G_2^* into H by the Fisher's construction (see e.g. book [10]) described next.

Definition 7. *Let G be a graph. Let $\sigma = (\sigma_v)_{v \in V(G)}$ be a choice, for every vertex v, of a linear ordering of the edges incident to v. The* blow-up, *or* Δ-extension, *of (G, σ) is the graph G^σ obtained by performing the following operation one by one for each vertex v. Let $e_1, \ldots, e_d$ be the linear ordering σ_v and let $e_i = vu_i$, $i = 1, \ldots, d$. We delete the vertex v and replace it with a path consisting of $6d$ new vertices $v_1, \ldots, v_{6d}$ and edges $v_i v_{i+1}$, $i = 1, \ldots, 6d - 1$. To this path, we add edges $v_{3j-2}v_{3j}$, $j = 1, \ldots, 2d$. Finally, we add edges $v_{6i-4}u_i$ corresponding to the original edges $e_1, \ldots, e_d$.*

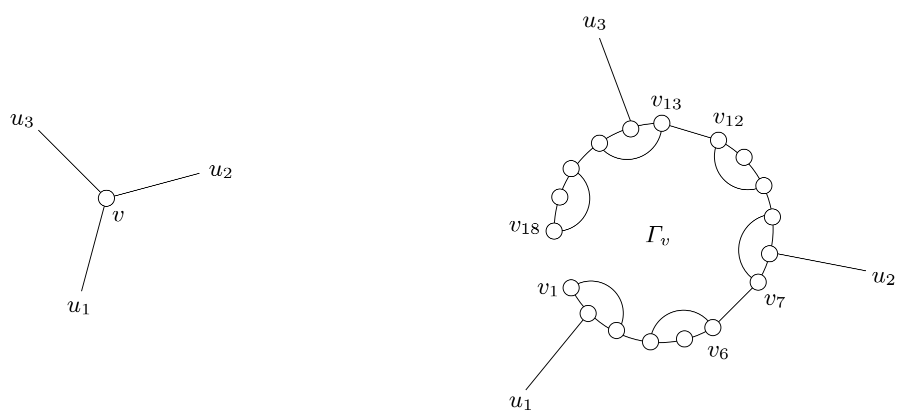

Fig. 5. For a node v with the neighborhood illustrated in 5(a) the associated gadget Γ_v is depicted in 5(b).

The subgraph of G^σ spanned by the $6d$ vertices $v_1, \ldots, v_d$ that replaced a vertex v of the original graph will be called a *gadget* and denoted by Γ_v. The edges of G^σ which do not belong to a gadget are in natural bijection with the edges of G. By abuse of notation, we will identify an edge of G with the corresponding edge of G^σ. Thus $E(G^\sigma)$ is the disjoint union of $E(G)$ and the various $E(\Gamma_v)$ $(v \in V(G))$.

It is important to note that different choices of linear orderings σ_v at the vertices of G may lead to non-isomorphic graphs G^σ. Nevertheless, one always has the following:

Lemma 1. *There is a natural bijection between the set of even subsets of G and the set of perfect matchings of G^σ. More precisely, every even set $E' \subseteq E(G)$*

uniquely extends to a perfect matching $M \subset E(G^\sigma)$, and every perfect matching of G^σ arises (exactly once) in this way.

It follows that if we set the weights of the edges of the gadgets of $(G_2^*)^\sigma$ equal to zero, we get that the value of the MAX-CUT problem for G is equal to the max REM of $H = (G_2^*)^\sigma$. This finishes the proof of Theorem 2.

References

1. S. Bravyi, Contraction of matchgate tensor networks on non-planar graphs, Contemporary Mathematics 482 (2009).
2. D. Cimasoni, N. Reshetikhin, Dimers on surface graphs and spin structures I. Comm. Math.Phys. 275(1) (2007).
3. R. Curticapean, M. Xia, Parametrizing the Permanent: Genus, Apices, Evaluation mod 2^k, Foundations of Computer Science (FOCS), IEEE Computer Society (2015).
4. A. Galluccio, M. loebl, On the theory of pfaffian orientations II. t-joins, k-cuts, and duality of enumeration, Electron. J. Combin. 6 (1999).
5. A. Galluccio, M. Loebl, J. Vondrak, A new algorithm for the Ising problem: partition function for finite lattice graphs, Physical Review Letters 84 (2000).
6. A. Galluccio, M. Loebl, J. Vondrak, Optimization via enumeration: a new algorithm for the Max Cut problem, Mathematical Programming 90 (2001).
7. R. Impagliazzo, R. Paturi, On the Complexity of k-SAT, J. Comput. Syst. Sci., 62(2) (2001), 367-375.
8. R. Impagliazzo, R. Paturi, F. Zane, Which Problems Have Strongly Exponential Complexity?, J. Comput. Syst. Sci., 63(4) (2001), 512-530.
9. P. Kasteleyn, Graph theory and crystal physics, in: Graph theory and theoretical physics, Academic Press, London (1967).
10. M. Loebl, Discrete mathematics in statistical physics, Vieweg and Teubner 2010.
11. M. Loebl, G. Masbaum, On the optimality of the Arf invariant formula for graph polynomials, Advances in Mathematics 226 (2011).
12. J. Lukic, A. Galluccio, E. Marinari, O.C. Martin, G. Rinaldi, Critical Thermodynamics of the Two-Dimensional +/-J Ising Spin Glass, Physical Review Letters 92 (11) (2004).
13. K. Mulmuley, U.V. Vazirani, V.V. Vazirani, Matching is as easy as matrix inversion, In Proceedings of the 19th Annual ACM Symposium on Theory of Computing, STOC'87 (1987), 345– 354.
14. J. Schaefer, The complexity of satisfiability problems, 10th Annual ACM Symposium on Theory of Computing (STOC) 216- 226 (1978).
15. G. Tesler, Matchings in graphs on Non-orientable Surfaces, Journal of Combinatorial Theory-B, 78 (2000).
16. L. Valiant, Holographic Algorithms (Extended Abstract), Foundations of Computer Science (FOCS), IEEE Computer Society (2004).
17. L. Valiant, The complexity of computing the permanent, Theoretical Computer Science 8,2 (1979).