



HAL
open science

Isoradial immersions

Cédric Boutillier, David Cimasoni, Béatrice de Tilière

► **To cite this version:**

Cédric Boutillier, David Cimasoni, Béatrice de Tilière. Isoradial immersions. *Journal of Graph Theory*, In press, 99 (4), pp.715-757. 10.1002/jgt.22761 . hal-02423791

HAL Id: hal-02423791

<https://hal.science/hal-02423791v1>

Submitted on 4 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright

Isoradial immersions

Cédric Boutillier*, David Cimasoni^{†‡}, Béatrice de Tilière[§]

Abstract

Isoradial embeddings of planar graphs play a crucial role in the study of several models of statistical mechanics, such as the Ising and dimer models. Kenyon and Schlenker [KS05] give a combinatorial characterization of planar graphs admitting an isoradial embedding, and describe the space of such embeddings. In this paper we prove two results of the same type for generalizations of isoradial embeddings: *isoradial immersions* and *minimal immersions*. We show that a planar graph admits a flat isoradial immersion if and only if its train-tracks do not form closed loops, and that a bipartite graph has a minimal immersion if and only if it is minimal. In both cases we describe the space of such immersions. The techniques used are different in both settings, and distinct from those of [KS05]. We also give an application of our results to the dimer model defined on bipartite graphs admitting minimal immersions.

Keywords— minimal graph, bipartite graph, isoradial embedding, minimal immersion, isoradial immersion

1 Introduction

An *isoradial embedding* of a planar graph G is an embedding of G into the plane so that each face is inscribed in a circle of radius 1, and so that the circumcenter of any given face is in the interior of that face. Joining each vertex of G with the circumcenters of the two adjacent faces, we see that such an isoradial embedding is equivalent to a so-called *rhombic embedding* of the associated *quad-graph* G^\diamond , *i.e.*, an embedding where each face of G^\diamond is given by a rhombus of side-length 1. This is illustrated in Figure 1.

*Sorbonne Université, CNRS, Laboratoire de Probabilités Statistique et Modélisation, LPSM, UMR 8001, F-75005 Paris, France; Institut Universitaire de France. cedric.boutillier@sorbonne-universite.fr

[†]Université de Genève, Section de Mathématiques, 1211 Genève 4, Suisse. david.cimasoni@unige.ch

[‡]corresponding author

[§]PSL University-Dauphine, CNRS, UMR 7534, CEREMADE, 75016 Paris, France; Institut Universitaire de France. detiliere@ceremade.dauphine.fr




Figure 1: A piece of an isoradial embedding of an infinite planar graph G (in black) and the underlying quad-graph G^\diamond (in gray).

Rhombic embeddings were introduced by Duffin [Duf68], as a natural class of graphs for which the Cauchy–Riemann equations admit a nice discretization. These ideas were rediscovered by Mercat [Mer01], who built an extensive theory of discrete holomorphicity in one complex variable. Kenyon [Ken02] introduced the concept of isoradial embedding as defined above, and proved it to be very useful in the study of the dimer model and the Green function [Ken02]. Several other models of statistical mechanics, such as the Ising model, are also shown to exhibit remarkable properties when studied on isoradial graphs with well-chosen weights, see e.g. [BdT10, BdT11, Cim12]. One of the deep underlying reasons lies in the star-triangle transformation and the corresponding Yang–Baxter equations, whose solutions admit an explicit parameterization when the graph of the model is isoradially embedded, see [Bax89] and references therein. This culminates in the work of Chelkak–Smirnov who developed the theory of discrete complex analysis on isoradial graphs to a high point of sophistication [CS11], and used it to prove conformal invariance and universality of the Ising model [CS12] precisely on this class of graphs. The study of statistical mechanics on isoradial graphs is still a very active area of research, as shown by the recent results obtained on bond percolation [GM14, Gri14], dimer heights [dT07, Li17], the random cluster model [BDCS15, DCLM18], random rooted spanning forests [BdTR17], and the Ising model [BdTR19]. We refer to the survey [BdT12] for more details on the history of isoradial graphs and their role in statistical physics.

Coming back to the strictly combinatorial study of isoradial graphs, a natural question immediately arises: when does a fixed planar graph G admit an isoradial embedding,

and if so, what is the space of such embeddings of G ? Let us briefly explain the answer to this question obtained by Kenyon and Schlenker in [KS05]. Given an arbitrary planar, embedded graph G , replace each edge by two (non-oriented) segments crossing each other at this edge, as described in Figure 4. Gluing these segments together, we obtain a set of planar curves called *train-tracks*, see Section 2.1 for a more formal definition. The main result of [KS05] is that a graph G admits an isoradial embedding if and only if its train-tracks form neither closed loops nor self-intersections, and two train-tracks intersect at most once. To describe the space of isoradial embeddings of such a graph G , let us denote by Z_G the space of maps α associating to each oriented train-track \vec{t} of G a direction $\alpha(\vec{t}) \in \mathbb{R}/2\pi\mathbb{Z}$, so that the same train-track with the opposite orientation \vec{t} receives the opposite direction $\alpha(\vec{t}) = \alpha(\vec{t}) + \pi$. An isoradial embedding of G defines an element of Z_G as follows: assign to an oriented train-track the direction of the parallel sides of the rhombi crossing this train-track, say from left to right. Note that any two oriented train-tracks \vec{t}_1, \vec{t}_2 intersecting (once) define a rhombus, so the corresponding two directions must satisfy some condition for this rhombus to be embedded in the plane with the right orientation: with the orientations of Figure 4 and the conventions above, we need to have $\alpha(\vec{t}_2) - \alpha(\vec{t}_1) \in (0, \pi)$. These conditions define a subspace of Z_G , which is shown in [KS05] to be equal to the space of isoradial embeddings of G .

The aim of the present paper is to prove two results of the same flavor for the largest possible families of planar graphs. We now state these two results, referring to Sections 3 and 4 for complete and precise statements.

As mentioned above, in the definition of an isoradial embedding, each rhombus has a *rhombus angle* $\theta = \alpha(\vec{t}_2) - \alpha(\vec{t}_1)$ in $(0, \pi)$. What if we relax this condition and allow rhombus angles in $(0, 2\pi)$, thus yielding folded rhombi? One of the goals of this paper is to prove that, in the bipartite case, this situation can be precisely understood. In the bipartite case indeed, it is possible to orient the train-tracks in a consistent way, *i.e.*, so that train-tracks turn counterclockwise around vertices of a given class of the bipartition and clockwise around vertices of the other class; such a consistent orientation is uniquely defined from the bipartite structure, and canonical up to complete inversion. Then allowing the rhombus angles to be in $(0, 2\pi)$ and asking that the total angle at vertices and faces is equal to 2π yields the notion of *minimal immersion*, see Sections 3.1 and 4.1 for a formal definition. The main result of Section 4 can be loosely stated as follows, see Theorem 23 in Section 4.1 for the precise statement.

Theorem 1. *Consider a bipartite, planar graph G with the consistent orientation of the train-tracks. Then, the graph G admits a minimal immersion if and only if its train-tracks do not self-intersect and two oriented train-tracks never intersect twice in the same direction. In any case, the space of minimal immersions of G is given by an explicit subspace Y_G of Z_G defined by one condition for each vertex and for each face of G .*

To be slightly more precise, if G satisfies the two conditions of Theorem 1, then the space Y_G is shown to contain (and in the \mathbb{Z}^2 -periodic case, to coincide with) the non-

empty subspace X_G of angle maps which are monotone with respect to a natural cyclic order on the set of oriented train-tracks of G , see Section 4.1.

Graphs satisfying these two conditions are referred to as *minimal graphs*. They were introduced by Thurston [Thu17] (the original arXiv paper appeared in 2004), play an important role in the string theory community [Gul08, IU11] where they are referred to as *consistent dimer models*, and in the seminal paper of Goncharov–Kenyon [GK13]. In the specific case of \mathbb{Z}^2 -periodic graphs, some of the concepts we use already appear in the literature: local and global ordering around vertices [IU11, Lemma 4.2], minimal immersions (in an understated way) in the proof of [GK13, Lemma 3.9], the space X_G [GK13, Theorem 3.3].

As explained above, the study of isoradial embeddings is motivated by statistical physics, and the present work is no exception. Indeed, in Theorem 31, we use the result above to show that there exists $\alpha \in Z_G$ such that some naturally occurring associated edge-weights [Ken02, Foc15] satisfy the so-called *Kasteleyn condition* if and only if the graph G is minimal, see Section 4.3. This allows us to harness Kasteleyn theory [Kas61, Kup98] in our papers [BCdT20, BCdT21], which deal with the dimer model on minimal graphs with the weights of [Foc15]. As a consequence, minimal graphs form the largest class of graphs where dimer models with such weights can be solved in a satisfactory way. Otherwise said, going beyond minimal immersions would give models with negative weights whose probabilistic interpretation is problematic. Moreover, the space Y_G of minimal immersions of a given minimal graph G coincides with the space of dimer models on G with fixed “abstract spectral data” (*i.e.* an M -curve together with a divisor on it, see [BCdT21]). Also, the minimal immersion of G corresponding to an element $\alpha \in Y_G$ gives a natural embedding of G^* , as it can be shown to coincide in the rational limit with the t -embedding of G^* described in [KLRR18], see [BCdT20, Section 8.1]. Finally, it is likely that a significant amount of the discrete complex analysis of [CS11] extends from isoradial to minimal graphs in a rather straightforward way, but we did not explore this question. Summarizing, it seems that minimal graphs rather than isoradial ones give the most natural framework to study such aspects of statistical mechanics models.

Although in the bipartite case, minimal graphs form a larger family than isoradial ones, they are still quite specific amongst planar graphs. With the goal of “immersing” with rhombi the most general possible planar graphs, we remove the assumption of being bipartite (as for isoradial embeddings) and we furthermore relax the conditions on the rhombus angles: in Section 3, we allow for the rhombus angle θ to be any lift in \mathbb{R} of the angle $\alpha(\vec{t}_2) - \alpha(\vec{t}_1) \in \mathbb{R}/2\pi\mathbb{Z}$, and define the dual rhombus angle to be $\theta^* = \pi - \theta$. The resulting theory is trivial without additional condition, so we impose the following geometrically natural one: these angles add up to 2π around each vertex and face of G when pasting these “generalized rhombi” together using the combinatorial information of G . This leads to the notion of *flat isoradial immersion* of the planar graph G , whose formal definition can be found in Section 3.1. In order to keep track of the different notions introduced, let us emphasize that a minimal immersion is a flat isoradial immersion of a bipartite graph with all rhombus angles restricted to be

in $(0, 2\pi)$. There is a natural equivalence relation on flat isoradial immersions of a fixed graph, see Definition 14 in Section 3.2. We now state the main theorem of Section 3.

Theorem 2. *A planar graph G admits a flat isoradial immersion if and only if its train-tracks do not form closed loops. When this is the case, the space of equivalence classes of flat isoradial immersions of G is given by the space Z_G .*

One of the main tools in the proof of this theorem is a combinatorial result, Lemma 4 of Section 2.2 which, although slightly technical, might be of independent interest. Proving Theorem 2 amounts to solving a linear system defined on an auxiliary graph. In the specific \mathbb{Z}^2 -periodic case, it turns out that the same linear system appears in a completely different context, see the physics papers [FHV⁺06, Equations (2.2) and (2.3)] and [Gul08, Equations (3.1) and (3.2)] for R -charges in some supersymmetric conformal field theory. Therefore, even though the possible role of flat isoradial immersions in statistical mechanics and discrete complex analysis remains elusive, their intriguing appearance in superconformal field theory does suggest interesting further research.

We conclude this introduction by pointing out that, even though Theorem 1, Theorem 2 and the main result of [KS05] are of the same nature, they belong to different categories: isoradial embeddings are purely geometric objects while flat isoradial immersions are almost entirely combinatorial, and minimal immersions lie in between. As a consequence, the sets of tools used in the proofs of these three statements are almost entirely disjoint. For example, the proof of the “only if” direction is immediate in [KS05, Theorem 3.1], but requires a discrete Gauss–Bonnet formula in Theorem 1, and the combinatorial study of train-track deformations in Theorem 2. We refer the reader to Figure 2 for a visualization of the classes of graphs and of the geometric realizations involved in these three results.

This paper is organised as follows.

- Section 2 deals with the purely combinatorial aspects of our work: the fundamental notions are defined in Section 2.1, while Section 2.2 contains a crucial result on train-tracks of an arbitrary planar graph. Section 2.3 deals with oriented train-tracks of bipartite graphs, defining a natural cyclic order on them, and studying this order in the case of minimal graphs.
- In Section 3, we first introduce the general notion of (flat) isoradial immersion (Section 3.1), then state the precise version of Theorem 2 in Section 3.2, and give the proof in Section 3.3. The last Section 3.4 gathers additional remarks on the definition of isoradial immersions, and on the proof of Theorem 2. In particular, we provide a geometric interpretation of isoradial immersions using rhombi that are iteratively folded along their primal or dual edge; we also prove a discrete Gauss–Bonnet formula.
- Section 4 deals with Theorem 1 above. In Section 4.1, we define minimal immersions of bipartite graphs, give a geometric characterization of these immersions,

define the spaces X_G and Y_G and state the main result. All proofs are contained in Section 4.2, while Section 4.3 deals with the aforementioned implications on the dimer model.

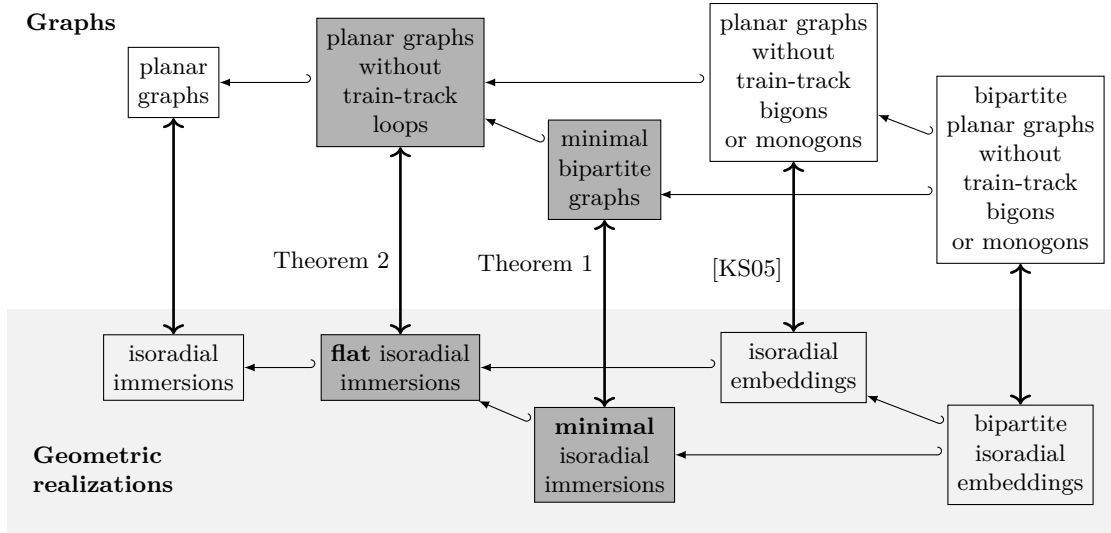


Figure 2: Visualization of Theorems 1 and 2 (in shaded boxes) and the main result of [KS05]. Inside each horizontal half, the arrows connecting the various classes of planar graphs (resp. geometric realizations) indicate inclusions. The vertical double arrows indicate an equivalence between a graph belonging to a particular class and it admitting a particular geometric realization. The first vertical arrow stems from the definition of isoradial immersions (see Section 3.1), whereas the last one is a specialization of the results of [KS05] to the bipartite case.

Acknowledgments

This project was started when the second-named author was visiting the first and third-named authors at the LPSM, Sorbonne Université, whose hospitality is thankfully acknowledged. The authors would also like to thank the anonymous referees for their constructive remarks. The first- and third-named authors are partially supported by the *DIMERS* project ANR-18-CE40-0033 funded by the French National Research Agency. The second-named author is partially supported by the Swiss National Science Foundation.

2 Planar graphs and train-tracks

The aim of this section is to introduce the basic combinatorial notions that are used throughout this article. We also prove four combinatorial lemmas that play a crucial

role in the proof of the main results of Sections 3 and 4.

2.1 General definitions

Consider a locally finite graph $G = (V, E)$ embedded in the plane so that its bounded faces are topological discs, and denote by $G^* = (V^*, E^*)$ the dual embedded graph. To avoid immediately diving into technical details, we first assume that all faces of G are bounded; we then turn to the general case, where unbounded faces require extra care.

Given a graph G as above, the associated *quad-graph* G^\diamond has vertex set $V \sqcup V^*$ and edge set defined as follows: a primal vertex $v \in V$ and a dual vertex $f \in V^*$ are joined by an edge each time the vertex v lies on the boundary of the face corresponding to the dual vertex f . Note that a vertex can appear twice on the boundary of a face, which gives rise to two edges between the corresponding vertices of G^\diamond . Since the faces of G are topological discs, the quad-graph G^\diamond embeds in the plane with faces consisting of (possibly degenerate) quadrilaterals. For example, a degree 1 vertex or a loop in G gives rise to a quadrilateral with two adjacent sides identified, see Figure 3 (left).

Following [KS05], see also [Ken02], we define a *train-track* as a maximal path of adjacent quadrilaterals of G^\diamond which does not turn: when it enters a quadrilateral, it exits through the opposite edge. Train-tracks are also known as rapidity lines in the field of integrable systems, see for example [Bax78]. By a slight abuse of terminology, we actually think of a train-track as the corresponding path in the dual graph $(G^\diamond)^*$ crossing opposite sides of the quadrilaterals. (These notions still make sense when a quadrilateral is degenerate.) Note that the graphs G and G^* have the same set of train-tracks, which we denote by \mathcal{T} . Note also that the local finiteness of G implies that \mathcal{T} is at most denumerable. We define the *graph of train-tracks* of G , denoted by $G^\mathcal{T}$, as the planar embedded graph $(G^\diamond)^*$.

By construction, its set of vertices $V^\mathcal{T}$ corresponds to the edges of G (hence the forthcoming notation $e \in V^\mathcal{T}$), while its faces correspond to vertices and faces of G . By definition, the graph G and its dual G^* define the same graph of train-tracks $G^\mathcal{T} = (G^*)^\mathcal{T}$, which is 4-regular: a typical edge $v_1 v_2 \in E$ and corresponding train-tracks $t_1, t_2 \in \mathcal{T}$ are illustrated in Figure 4.

We now turn to the general case, where G can admit one or several unbounded faces. In such a case, first construct the associated graph $(G^\diamond)^*$ as above, considering each unbounded face as a vertex of the dual graph G^* . Then, define $G^\mathcal{T}$ as the subgraph of $(G^\diamond)^*$ obtained by removing all edges included in an unbounded face of G ; finally, define \mathcal{T} as the set of maximal paths in $G^\mathcal{T}$ crossing opposite sides of the quadrilaterals. (Note that the embedding of $(G^\diamond)^*$ into the plane might depend on the embedding of G^* , but the planar embedded graph $G^\mathcal{T} \subset \mathbb{R}^2$ is fully determined by $G \subset \mathbb{R}^2$.) By construction, its set of vertices $V^\mathcal{T}$ corresponds to the edges of G , while its set of faces $F^\mathcal{T}$ corresponds to *inner vertices*, not adjacent to an unbounded face, and (bounded) *inner faces* of G . Note that in general, the graph $G^\mathcal{T}$ is no longer 4-regular: indeed, unbounded faces give rise to vertices of degree 3 and 2, where one train-track, or both, stop in the middle of the corresponding quadrilateral. Note finally that in the presence of

unbounded faces, the graph G and its dual G^* no longer define the same graph of train-tracks. However, this notion remains (locally) self-dual away from unbounded faces. We refer to Figure 3 for an illustration of a graph of train-tracks with all the possible pathologies.




Figure 3: Left: a planar graph G (black). The dual G^* has vertices corresponding to bounded faces, marked with grey diamonds, and an additional vertex for the unbounded face of G , which is here at infinity. The quad-graph G° is represented with grey edges. Those with a dashed end are connected to the dual vertex at infinity. Train-tracks of G , paths of adjacent quadrilaterals, are materialized by colored lines (a : blue, b : green, c : orange, d : cyan). Right: the corresponding graph of train-tracks G^T .

We now come to orientation of train-tracks. Each train-track $t \in \mathcal{T}$ can be oriented in two directions: we let $\vec{t}, \overleftarrow{t}$ denote its oriented versions, and $\overline{\mathcal{T}}$ be the set of such oriented train-tracks. Given an inner oriented edge (v_1, v_2) of G and an oriented train-track \vec{t} crossing it, we say that these orientations are *coherent* if (v_1, v_2) and \vec{t} define the positive (counterclockwise) orientation of the plane, see Figure 4. We moreover fix the subscripts 1, 2 so that \vec{t}_1, \vec{t}_2 also defines the positive orientation of the plane.




Figure 4: An edge v_1v_2 of G and the two corresponding train-tracks oriented so that (v_1, v_2) and \vec{t}_1 , resp. \vec{t}_2 , are coherent, and labeled so that \vec{t}_1, \vec{t}_2 defines a positive orientation of the plane.

In the case of bipartite graphs, there exists a natural orientation of train-tracks, defined as follows. Recall that a graph G is *bipartite* if its vertex set can be partitioned into two

disjoint sets $V = B \sqcup W$ of *black* and *white* vertices, such that each edge of G joins a black vertex b to a white one w . Then, a *consistent* orientation of the train-tracks is obtained by requiring train-tracks to turn counterclockwise around white vertices and clockwise around black ones. By construction, this is the unique orientation which is coherent at every edge with the orientation of the edge from the white to the black vertex.

Without any hypothesis on G , the associated train-tracks can display all sorts of behavior: for example, a train-track might form a closed loop, it might intersect itself, and two train-tracks might intersect more than once. As mentioned in the introduction, the main result of [KS05] is that a graph G admits an isoradial embedding if and only if its train-tracks never behave in this way. The main theorem of Section 3 should be seen as a result of the same type: a graph G admits a more general geometric realization, called *isoradial immersion*, if and only if its train-tracks do not form closed loops. Finally, the main theorem of Section 4 is of the same flavor: a bipartite graph G with consistently oriented train-tracks admits a so-called *minimal immersion* if and only if its oriented train-tracks do not form closed loops, self-intersections, or parallel bigons, where a *parallel bigon* is a pair of oriented train-tracks that intersect twice at two points x and y and both are oriented from x to y [Thu17]. The three forbidden configurations are illustrated in Figure 5.




Figure 5: Train-tracks form a closed loop (left), a self-intersection (middle), or a parallel bigon (right), if G^J contains the corresponding graph as a minor.

2.2 Cycles in the graph of train-tracks

The basic definitions being given, we now turn to the first result of this article, Lemma 4, a fairly technical combinatorial statement. This lemma is used in Section 3 to describe the space of solutions of a linear system indexed by vertices of G^\diamond . But we believe that under its general form here, it is interesting in itself and may have applications to other contexts.

Following the standard terminology, we say that a finite subgraph of an arbitrary abstract graph Γ is a *cycle*, resp. a *simple cycle*, if all of its vertices are of even degree, resp. of degree 2. We denote by $\mathcal{E}(\Gamma)$ the set of cycles of Γ (\mathcal{E} stands for *even*). Recall that the symmetric difference endows $\mathcal{E}(\Gamma)$ with the structure of an \mathbb{F}_2 -vector space.

An intermediate step for Lemma 4 consists in exhibiting a bijection between the cycles of

$G^{\mathcal{J}}$ and those of a natural auxiliary graph. Before establishing this bijection in Lemma 3, we discuss some features of cycles of $G^{\mathcal{J}}$.

Vertices of a cycle $c \in \mathcal{E}(G^{\mathcal{J}})$ fall into three categories: degree 4 vertices where by definition, c locally follows two (non-necessarily distinct) intersecting train-tracks; degree 2 vertices where c stays on the same train-track, *i.e.*, where it crosses opposite sides of the corresponding quadrilateral; and degree 2 vertices where c switches train-tracks, *i.e.*, where it “turns” inside the quadrilateral. We call these latter vertices *corners* of c .

From the above observations we define a procedure for drawing c on the graph, giving a unique decomposition of c into *closed curves*: start from an edge of c , follow the cycle in one of the two possible directions by going straight at each intersection. When coming back to the starting point, we have covered a cycle component of c in our decomposition, referred to as a *closed curve*. If the whole of c has been covered, stop; else remove this component from c , and iterate the procedure from an uncovered edge of c . Since the cycle is finite, the algorithm of course terminates. Note also that specifying the type of behavior at the degree 4 vertices yields uniqueness of the decomposition of c into closed curves, and that the behaviour chosen minimizes the number of corners: the total number of corners in all closed curves appearing in this decomposition of c is the same as the number of corners of c itself, see Figure 6 for an example.




Figure 6: Example of a cycle (left) with corners partitioned into two parts (with different shades of purple), corresponding to the corners of the closed curves (right) entering into its canonical decomposition into closed curves.

When none of the train-tracks of G form a closed loop, any closed curve has at least one corner. In this case, we relate cycles on $G^{\mathcal{J}}$ with cycles on another, auxiliary abstract graph $G' = (V', E')$, constructed from $G^{\mathcal{J}}$ by setting $V' = \mathcal{T}$ and $E' = V^{\mathcal{J}}$. More precisely, vertices of G' are given by the train-tracks, and each intersection in $G^{\mathcal{J}}$ between two train-tracks (resp. self-intersection of one train-track) defines an edge between the two corresponding vertices of G' (resp. a loop at the corresponding vertex of G'), see Figure 7.

Observe that any closed curve $c \subset G^{\mathcal{J}}$ defines a cycle $\varphi(c) \subset G'$ with the corners of c corresponding to edges of $\varphi(c)$. We use the unique decomposition of cycles into closed curves described above to extend φ on $\mathcal{E}(G^{\mathcal{J}})$ by declaring that φ sends a cycle to the sum of the images of its components in that particular decomposition.

Lemma 3. *Let $G^{\mathcal{J}}$ be a connected graph of train-tracks, none of which forms a closed loop. Then, the function $\varphi: \mathcal{E}(G^{\mathcal{J}}) \rightarrow \mathcal{E}(G')$ is an isomorphism of \mathbb{F}_2 -vector spaces.*




Figure 7: The auxiliary graph G' appearing in the proof of Lemma 4, for the example of graph G given in Figure 3.

Proof. First, φ is indeed a linear map. This can be checked by noticing that when decomposing a cycle c as a sum of smaller cycles $c = \sum_i c_i$, the cardinality of the set of indices i such that $e \in V^{\mathcal{J}}$ is a corner of c_i is odd if e is a corner of c , and even if not. Because we look at linear combinations with coefficients mod 2, this is enough to conclude that $\varphi(c) = \sum_i \varphi(c_i)$.

The application φ is clearly onto: all simple cycles of G' can be realized, and they generate the full space of cycles $\mathcal{E}(G')$. We now wish to show that φ is injective.

Let us first assume that $G^{\mathcal{J}}$ is finite. In such a case, the auxiliary graph G' is finite as well. Since we assume $G^{\mathcal{J}}$ to be connected, so is G' , and we can determine the dimension of $\mathcal{E}(G')$ using the standard computation

$$1 - \dim(\mathcal{E}(G')) = \chi(G') = |V'| - |E'| = |\mathcal{J}| - |V^{\mathcal{J}}|,$$

where χ stands for the Euler characteristic.

To evaluate the dimension of $\mathcal{E}(G^{\mathcal{J}})$, let us denote by $V^{\mathcal{J}} = V_2 \sqcup V_3 \sqcup V_4$ the partition of the vertices of $G^{\mathcal{J}}$ into vertices of degree 2, 3 and 4. We have the obvious equality $2|V_2| + 3|V_3| + 4|V_4| = 2|E^{\mathcal{J}}|$, and the slightly less obvious one $2|V_2| + |V_3| = 2|\mathcal{J}|$ which uses the fact that $G^{\mathcal{J}}$ has no closed loops. Together, they lead to the equation $2|V^{\mathcal{J}}| = |E^{\mathcal{J}}| + |\mathcal{J}|$, giving

$$1 - \dim(\mathcal{E}(G^{\mathcal{J}})) = \chi(G^{\mathcal{J}}) = |V^{\mathcal{J}}| - |E^{\mathcal{J}}| = |\mathcal{J}| - |V^{\mathcal{J}}|. \quad (1)$$

By the two equations displayed above, we see that $\varphi: \mathcal{E}(G^{\mathcal{J}}) \rightarrow \mathcal{E}(G')$ is a surjective linear map between finite dimensional vector spaces of the same dimension, and therefore an isomorphism.

The general (possibly infinite) case can be seen as a consequence of the finite case, as follows. Since $\varphi = \varphi_G: \mathcal{E}(G^{\mathcal{J}}) \rightarrow \mathcal{E}(G')$ is already known to be a surjective linear map, we are left with proving that the only $c \in \mathcal{E}(G^{\mathcal{J}})$ such that $\varphi(c)$ vanishes is $c = 0$, the empty cycle. Such a c is contained in a finite connected subgraph $H^{\mathcal{J}}$ of $G^{\mathcal{J}}$; let H' be the corresponding auxiliary graph, and $\varphi_H: \mathcal{E}(H^{\mathcal{J}}) \rightarrow \mathcal{E}(H')$ be the corresponding linear

map, which we know is an isomorphism since $\mathbf{H}^{\mathcal{J}}$ is finite, connected, and contains no closed loop. By naturality of φ , the following diagram commutes

$$\begin{array}{ccc} \mathcal{E}(\mathbf{G}^{\mathcal{J}}) & \xrightarrow{\varphi_{\mathbf{G}}} & \mathcal{E}(\mathbf{G}') \\ j \uparrow & & j' \uparrow \\ \mathcal{E}(\mathbf{H}^{\mathcal{J}}) & \xrightarrow{\varphi_{\mathbf{H}}} & \mathcal{E}(\mathbf{H}'), \end{array}$$

where the vertical arrows are the inclusions of spaces of cycles induced by the inclusion of subgraphs. As mentioned, there exists $c_{\mathbf{H}} \in \mathcal{E}(\mathbf{H}^{\mathcal{J}})$ such that $c = j(c_{\mathbf{H}})$, leading to

$$0 = \varphi_{\mathbf{G}}(c) = \varphi_{\mathbf{G}}(j(c_{\mathbf{H}})) = j'(\varphi_{\mathbf{H}}(c_{\mathbf{H}})).$$

Since $j' \circ \varphi_{\mathbf{H}}$ is injective, it follows that $c_{\mathbf{H}}$ is equal to 0, and so is $c = j(c_{\mathbf{H}})$. \square

Using the isomorphism φ to transport a natural basis of $\mathcal{E}(\mathbf{G}')$ to $\mathcal{E}(\mathbf{G}^{\mathcal{J}})$ yields the following combinatorial statement about $\mathbf{G}^{\mathcal{J}}$.

Lemma 4. *Let $\mathbf{G}^{\mathcal{J}}$ be a connected graph of train-tracks, none of which forms a closed loop. Then, there exists a set $M \subset \mathbf{V}^{\mathcal{J}}$, an injective map $\tau: M \rightarrow \mathcal{T}$, and a basis $(c_e)_{e \in \mathbf{V}^{\mathcal{J}} \setminus M}$ of $\mathcal{E}(\mathbf{G}^{\mathcal{J}})$ consisting of closed curves indexed by the vertices of $\mathbf{V}^{\mathcal{J}} \setminus M$, such that:*

1. *each vertex $e \in M$ belongs to the associated train-track $\tau(e) \in \mathcal{T}$, and the image of τ consists of all train-tracks but one;*
2. *for every $e \in \mathbf{V}^{\mathcal{J}} \setminus M$, the vertex e is a corner of c_e , and all other corners of c_e belong to M .*

Recall that we use the notation e for vertices of $\mathbf{G}^{\mathcal{J}}$ because they correspond to edges of \mathbf{G} . This statement is quite technical, but its translation in plain English less so, at least in the finite case: there exists a set M of $|\mathcal{T}| - 1$ “marked vertices” (M stands for “marked”) such that each vertex e in this subset has a distinct train-track $\tau(e)$ associated by the map τ ; and a basis $(c_e)_{e \in \mathbf{V}^{\mathcal{J}} \setminus M}$ of cycles of $\mathbf{G}^{\mathcal{J}}$ consisting of closed curves, such that one can mark all vertices of $\mathbf{G}^{\mathcal{J}}$ starting from M and then marking the unique missing corner of each closed curve c_e . Note that this statement is easily seen not to hold when some train-track makes a loop.

Proof of Lemma 4. Since $\mathbf{G}^{\mathcal{J}}$ is connected, so is \mathbf{G}' which therefore contains a spanning tree \mathbf{T}' , rooted at an arbitrary vertex $t_0 \in \mathbf{V}' = \mathcal{T}$. Let $M \subset \mathbf{V}^{\mathcal{J}} = \mathbf{E}'$ denote the set of vertices of $\mathbf{G}^{\mathcal{J}}$ corresponding to the edges of \mathbf{T}' , and $\tau: M \rightarrow \mathcal{T}$ be the map defined by associating to each edge of \mathbf{T}' its adjacent vertex in the direction opposite to the root t_0 . Since \mathbf{T}' has no cycle, the map τ is injective, and since \mathbf{T}' spans all vertices, the image of τ is $\mathcal{T} \setminus \{t_0\}$. This shows the first point.

To conclude the proof, write ε_e for the element of $\mathbf{E}' \setminus \mathbf{T}'$ corresponding to $e \in \mathbf{V}^{\mathcal{J}} \setminus M$, let $\varepsilon'_e \in \mathcal{E}(\mathbf{G}')$ be the unique (simple) cycle in $\mathbf{T} \cup \{\varepsilon_e\}$, and set $c_e = \varphi^{-1}(\varepsilon'_e) \in \mathcal{E}(\mathbf{G}^{\mathcal{J}})$. The

crucial fact is that c_e is a closed curve. Indeed, ε'_e is a simple cycle. Every pair of edges of ε'_e connected to the same vertex t correspond to subsequent corners of c_e connected by the train-track t . Therefore, when performing the decomposition procedure into closed curves on the cycle c_e , we obtain a unique component, implying that c_e is indeed a closed curve.

Since $(\varepsilon'_e)_{e \in V^{\mathcal{T}} \setminus M}$ is a basis of $\mathcal{E}(\mathbf{G}')$ and φ an isomorphism, $(c_e)_{e \in V^{\mathcal{T}} \setminus M}$ is a basis of $\mathcal{E}(\mathbf{G}^{\mathcal{T}})$. Via the isomorphism φ , the statement of the second point translates into the following tautology: for all $e \in V^{\mathcal{T}} \setminus M$, the edge ε_e belongs to ε'_e and all other edges of ε'_e belong to \mathbf{T}' . \square

2.3 Train-tracks in minimal bipartite graphs

In the whole of this section, we suppose that the planar embedded graph \mathbf{G} is bipartite, and has no unbounded faces. Recall from Section 2.1 that the train-tracks of \mathbf{G} can be consistently oriented so that black vertices are on the right, and white vertices on the left of the directed paths. Recall also that the above orientation is coherent at every edge with the orientation of train-tracks arising from oriented edges, directed from white vertices to black ones. We denote by $\vec{\mathcal{T}}$ the corresponding set of oriented train-tracks, which is nothing but a copy of \mathcal{T} , with all train-tracks consistently oriented.

In most of the present section and Section 4, we focus on a special class of planar bipartite graphs known as *minimal graphs*, introduced by Thurston [Thu17]. These graphs are central in the seminal paper [GK13] by Goncharov–Kenyon; they also occur in string theory where they are referred to as *consistent dimer models*, see [Gul08, IU11].

Definition 5. A planar, embedded, bipartite graph \mathbf{G} is said to be *minimal* if $\mathbf{G}^{\mathcal{T}}$ contains neither self-intersecting train-tracks, nor parallel bigons, see Figure 5.

Throughout the rest of this section, we assume that the train-tracks of \mathbf{G} do not form closed loops, a condition which always holds for minimal graphs. Indeed, with our conventions, no connected component of $\mathbf{G}^{\mathcal{T}}$ consists of a simple closed curve; therefore, a closed loop will either self-intersect, or meet another train-track and thus form a parallel bigon. Our goal is to prove two combinatorial results on the order of train-tracks of minimal bipartite graphs.

When the graph \mathbf{G} is \mathbb{Z}^2 -periodic, a partial cyclic order on train-tracks is defined as follows: the projection of train-tracks on the quotient graph \mathbf{G}/\mathbb{Z}^2 drawn on a torus are non-trivial oriented loops whose homology class in $H_1(\mathbb{T}; \mathbb{Z}) \simeq \mathbb{Z}^2$ is given by a pair of coprimes integers. The total cyclic order on such pairs thus induces a natural (partial) cyclic order on train-tracks, which has been exploited in the papers [Gul08, IU11, GK13].

We now extend this partial cyclic order to train-tracks of arbitrary bipartite planar graphs with neither train-track loops nor unbounded faces. Note that these conditions are equivalent to requiring that all train-tracks are bi-infinite (oriented) curves.

First, we say that two oriented train-tracks (or more generally, two non-closed oriented planar curves) are *parallel* if they satisfy one of the following conditions:

- they intersect infinitely many times in the same direction, see Figure 8, left,
- they are disjoint, and there exists a topological disc $D \subset \mathbb{R}^2$ that they cross in the same direction, see Figure 8, right.

Two oriented train-tracks \vec{t}_1, \vec{t}_2 are called *anti-parallel* if \vec{t}_1 and \vec{t}_2 are parallel.




Figure 8: Parallel train-tracks.

Let us now consider a triple of oriented train-tracks $(\vec{t}_1, \vec{t}_2, \vec{t}_3)$ of \mathbf{G} , pairwise non-parallel. Consider a compact disk B outside of which these train-tracks do not meet, apart from possible anti-parallel ones, and order $(\vec{t}_1, \vec{t}_2, \vec{t}_3)$ cyclically according to the outgoing points of the corresponding oriented curves in the circle ∂B . Note that since the three oriented train-tracks are pairwise non-parallel, this cyclic order does not depend on the choice of the disk B .

Definition 6. We call this partial cyclic order on $\vec{\mathcal{T}}$ the *global cyclic order*.

Remark 7.

- When \mathbf{G} is \mathbb{Z}^2 -periodic, it is easy to see that this notion of partial order coincides with the one described above. In this context, being parallel (resp. anti-parallel) is equivalent to having the same homology class (resp. opposite homology classes).
- In [KS05], Kenyon and Schlenker consider a (non-necessarily bipartite) planar graph \mathbf{G} whose train-tracks do not form closed loops, do not self-intersect, and such that any two train-tracks intersect at most once. The authors construct a topological embedding of $\mathbf{G}^{\vec{\mathcal{T}}}$ into the unit disc such that the image of each train-track is a smooth path connecting distinct boundary points of the disc, see [KS05, Lemma 3.2]. This construction can be adapted to our setting, yielding a continuous map from $\mathbf{G}^{\vec{\mathcal{T}}}$ to the unit disc, where parallel train-tracks connect the same boundary points of the disc. Furthermore, the (partial) cyclic order on $\vec{\mathcal{T}}$ given by the outgoing points of oriented train-tracks in the unit circle coincides with the global cyclic order defined above.

The first result is the elementary, yet fundamental observation that for minimal graphs, this global cyclic order on $\vec{\mathcal{T}}$ induces the standard cyclic order around any vertex. More formally, consider a bipartite planar graph \mathbf{G} , and a vertex $v \in \mathbf{V}$ of degree $n \geq 1$. The n incident edges are crossed by a set $\vec{\mathcal{T}}(v)$ of n oriented train-tracks strands, each one joining two consecutive edges, see Figure 9 (left). In general, it can happen that two of these strands belong to the same train-track, e.g. in the case of multiple edges. Furthermore, there is no reason in general for the global cyclic order on $\vec{\mathcal{T}}$ to restrict to the *local cyclic order* on $\vec{\mathcal{T}}(v)$, *i.e.*, the total cyclic order given by outgoing points of these strands. This actually holds if \mathbf{G} is minimal as stated by the following lemma. Note that in the \mathbb{Z}^2 -periodic case, this is essentially the content of Lemma 4.2. of [IU11].




Figure 9: Left: the set $\vec{\mathcal{T}}(v)$ of oriented train-tracks strands around the (white) vertex v . Right: Proof of Lemma 8.

Lemma 8. *Let \mathbf{G} be a planar, bipartite, minimal graph, and let v be an arbitrary vertex of \mathbf{G} . Then, the elements of $\vec{\mathcal{T}}(v)$ belong to distinct train-tracks, and the global cyclic order on $\vec{\mathcal{T}}$ restricts to the local cyclic order on $\vec{\mathcal{T}}(v)$.*

Proof. Fix a vertex v of a minimal bipartite graph \mathbf{G} , and consider the set $\vec{\mathcal{T}}(v)$ of adjacent train-tracks strands as illustrated in Figure 9 (left). First, observe that in order for two of these strands to belong to the same train-track, we need to connect the outgoing point of one of these strands with the ingoing point of another one. As one easily checks, this either introduces a self-intersection or a parallel bigon, thus contradicting minimality of \mathbf{G} . Finally, assume by means of contradiction that the global cyclic order on $\vec{\mathcal{T}}$ does not restrict to the local cyclic order on $\vec{\mathcal{T}}(v)$. Then, this inevitably creates a parallel bigon, see e.g. Figure 9 (right), contradicting the minimality of \mathbf{G} , and concluding the proof. \square

Our second result deals with the local order of train-tracks around faces. Since our graph \mathbf{G} is bipartite, any face f of \mathbf{G} is of even degree, say $2m$; in the degenerate case of a degree 1 vertex inside the face, the adjacent edge is counted twice.

The $2m$ edges bounding f are crossed by a set $\vec{\mathcal{T}}(f)$ of $2m$ oriented train-tracks strands, each one joining two consecutive boundary edges, see Figure 10. There is a natural partition $\vec{\mathcal{T}}(f) = \vec{\mathcal{T}}^\bullet(f) \sqcup \vec{\mathcal{T}}^\circ(f)$, where $\vec{\mathcal{T}}^\bullet(f)$ is the set of strands turning around a black

vertex of ∂f , *i.e.*, turning counterclockwise around f , and $\vec{\mathcal{T}}^\circ(f)$ is the set of strands turning around a white vertex of ∂f , *i.e.*, clockwise around f . In general, all the train-track strands in $\vec{\mathcal{T}}^\bullet(f)$ can belong to the same train-track, and similarly for $\vec{\mathcal{T}}^\circ(f)$: a face of degree 2 is the easiest example. Furthermore, there is no reason in general for the global cyclic order on $\vec{\mathcal{T}}$ to restrict to the *local cyclic orders* on $\vec{\mathcal{T}}^\bullet(f)$ and on $\vec{\mathcal{T}}^\circ(f)$, *i.e.*, the total cyclic orders given by outgoing points of these strands. Once again, this turns out to hold if G is minimal.




Figure 10: The set $\vec{\mathcal{T}}(f) = \vec{\mathcal{T}}^\circ(f) \sqcup \vec{\mathcal{T}}^\bullet(f)$.

Lemma 9. *Let G be a planar, bipartite, minimal graph, and let f be an arbitrary face of G . Then, there is a pair of train-track strands in $\vec{\mathcal{T}}^\bullet(f)$ which belong to distinct non-parallel train-tracks, and similarly for $\vec{\mathcal{T}}^\circ(f)$. Furthermore, the global cyclic order on $\vec{\mathcal{T}}$ restricts to the local cyclic orders on $\vec{\mathcal{T}}^\bullet(f)$ and on $\vec{\mathcal{T}}^\circ(f)$.*

Proof. Let us fix a face f of a bipartite minimal graph G , and assume by means of contradiction that all train-track strands in $\vec{\mathcal{T}}^\bullet(f)$ belong to the same train-track. Using the fact that train-tracks cannot self-intersect, we obtain that these strands are connected cyclically as illustrated in Figure 11 (left). This creates a parallel bigon, thus contradicting the minimality of G . Furthermore, two elements of $\vec{\mathcal{T}}^\bullet(f)$ cross the face f in anti-parallel fashion, so they cannot belong to parallel train-tracks (which have to be disjoint since the graph is minimal). The same argument holds for $\vec{\mathcal{T}}^\circ(f)$. Let us finally assume that the global cyclic order on $\vec{\mathcal{T}}$ does not restrict to the local cyclic orders on $\vec{\mathcal{T}}^\bullet(f)$ or on $\vec{\mathcal{T}}^\circ(f)$. In such a case, we inevitably have a self-intersection or a parallel bigon, see e.g. Figure 11 (right). This contradicts the minimality of G and concludes the proof. \square

A natural question is whether some converse to these lemmas holds.

Remark 10. In [IU11], the authors prove the following statement: if a \mathbb{Z}^2 -periodic bipartite planar graph G is such that its train-tracks do not form closed loops, do not




Figure 11: Proof of Lemma 9. Left: strands of $\vec{\mathcal{J}}^\bullet(f)$ coming from the same train-track. Right: global cyclic order not restricting to local cyclic orders of $\vec{\mathcal{J}}^\circ(f)$ and $\vec{\mathcal{J}}^\bullet(f)$. In both cases, shaded areas are bounded by parallel bigons or self-intersections.

self-intersect, two parallel train-tracks are always disjoint, and for any vertex v , the global cyclic order on $\vec{\mathcal{J}}$ restricts to the local cyclic order on $\vec{\mathcal{J}}(v)$, then G is minimal. Note however that the biperiodicity is used in a crucial way, more precisely the fact that the train-tracks define a finite number of asymptotic directions.

As a consequence of Theorem 23, we actually obtain a converse to these two lemmas: if a bipartite planar graph G has no train-track loops, no vertex of degree 1, and is such that the conclusions of Lemmas 8 and 9 hold, then G is minimal, see Corollary 28 of Section 4.

We conclude this section with the following:

Question. Given a minimal bipartite graph G , do the cyclic orders on $\vec{\mathcal{J}}(v)$ for all $v \in V$ and on $\vec{\mathcal{J}}^\bullet(f)$ and $\vec{\mathcal{J}}^\circ(f)$ for all $f \in F$ determine the global cyclic order on $\vec{\mathcal{J}}$?

As a consequence of our results, we will answer this question positively in the \mathbb{Z}^2 -periodic case, see Corollary 29 of Section 4.2.

3 Flat isoradial immersions of planar graphs

This section deals with a geometric notion, called *isoradial immersion* and defined in Section 3.1, which extends the classical notion of isoradial embedding. Our main result is stated in Section 3.2; it is the pendant in this more general context of Kenyon and Schlenker's theorem [KS05]; the proof is the subject of Section 3.3. One last Section 3.4 proves additional features of our theory.

3.1 Definitions

We start this section by defining the notion of *rhombic immersion*. Consider a planar graph G , the corresponding quad-graph G^\diamond , and the graph of train-tracks $G^{\mathcal{J}}$. Suppose that each pair of directed train-tracks $\vec{t}, \overleftarrow{t} \in \vec{\mathcal{J}}$ is assigned a pair of opposite directions $e^{i\alpha(\vec{t})}$ and $e^{i\alpha(\overleftarrow{t})} = -e^{i\alpha(\vec{t})}$; the angle $\alpha(\vec{t}), \alpha(\overleftarrow{t}) \in \mathbb{R}/2\pi\mathbb{Z}$ are referred to as the

angles of the train-tracks $\vec{t}, \overleftarrow{t}$, and we denote by Z_G the set of possible angle maps, i.e.

$$Z_G = \{ \alpha : \overleftarrow{\mathcal{J}} \rightarrow \mathbb{R}/2\pi\mathbb{Z} \mid \forall \vec{t} \in \overleftarrow{\mathcal{J}}, \alpha(\vec{t}) = \alpha(\overleftarrow{t}) + \pi \}.$$

Every element $\alpha \in Z_G$ defines an immersion of G^\diamond in \mathbb{R}^2 , by realizing every directed edge of G^\diamond crossed by a train-track \vec{t} from left to right as a translation of the unit vector $e^{i\alpha(\vec{t})}$. Every inner (quadrilateral) face of G^\diamond is mapped to a rhombus of unit-edge length in the plane. We call such an immersion a *rhombic immersion* of G^\diamond . Given the cyclic order around vertices in G^\diamond , the data of the rhombic immersion is equivalent to that of α .

Every face of G^\diamond corresponds to an edge e of G by construction. Using the notation of Section 2.1 and Figure 12, we associate to the oriented edge $e = (v_1, v_2)$, a *rhombus angle* $\theta_e \in [0, 2\pi)$, which is the unique lift in $[0, 2\pi)$ of $[\alpha(\vec{t}_2) - \alpha(\vec{t}_1)] \in \mathbb{R}/2\pi\mathbb{Z}$, where the square bracket denotes the equivalence class of a real number in $\mathbb{R}/2\pi\mathbb{Z}$. Recall from Section 2.1 that the orientation of \vec{t}_1 , resp. \vec{t}_2 , is chosen to be coherent with the orientation of e and that the subscripts 1, 2 are chosen so that \vec{t}_1, \vec{t}_2 defines a positive orientation of the plane. This rhombus angle is independent of the orientation of the edge e , as reversing the orientation of e also reverses the orientation of both train-tracks while keeping the subscripts fixed; for an internal edge e , it is the geometric angle of the rhombus image of the quadrilateral corresponding to e , at the vertices (corresponding to the images by the immersion of) v_1 and v_2 . The *dual rhombus angle*, measured at vertices f_1 or f_2 is given by $\theta_{e^*} = \pi - \theta_e$.




Figure 12: The image by a rhombic immersion of an inner (quadrilateral) face of G^\diamond , with vertices v_1 and v_2 corresponding to vertices of G , and f_1 and f_2 corresponding to faces of G . Edges are represented by unit vectors associated to train-tracks.

Note that in many papers, see e.g. [Ken02], the notation θ_e is used for the rhombus *half-angle*, but in our setting using θ_e for the rhombus *angle* is more convenient. There are three types of rhombi: *embedded* ones when θ_e belongs to $(0, \pi)$, *folded* ones for $\theta_e \in (\pi, 2\pi)$, and *degenerate* ones when $\theta_e \in \{0, \pi\}$.

To each edge e of G , we also assign an arbitrary lift $\tilde{\theta}_e \in \mathbb{R}$ of $[\alpha(\vec{t}_2) - \alpha(\vec{t}_1)] \in \mathbb{R}/2\pi\mathbb{Z}$. Since $\tilde{\theta}_e = \theta_e + 2\pi k_e$ for a unique $k_e \in \mathbb{Z}$, such a choice of lift amounts to the choice of an element $k = (k_e)_{e \in E} \in \mathbb{Z}^E$. In a similar way as above, to each edge e we associate a *generalized rhombus* with angle $\tilde{\theta}_e$ at the adjacent vertices, and dual angle $\tilde{\theta}_{e^*} = \pi - \tilde{\theta}_e$ at

the adjacent dual vertices. Geometrically, such a generalized rhombus should be thought of as being iteratively folded along its diagonals, see Section 3.4 for more details.

Definition 11. The *isoradial immersion* of \mathbf{G} defined by the train-track angles $\alpha \in Z_{\mathbf{G}}$ and the integers $k \in \mathbb{Z}^E$ is the gluing of the generalized rhombi using the combinatorial information of \mathbf{G} .

The surface obtained by gluing the (generalized) rhombi is naturally equipped with a flat metric with conical singularities at inner vertices and faces. On this surface, the graphs \mathbf{G}^\diamond and \mathbf{G} are naturally *embedded*. When projecting this surface onto the plane, the images of vertices and edges coincide with the images of the rhombic immersion defined by α . We define the cone angle at a vertex (resp. face) singularity as the sum of the lifts $\tilde{\theta}_e$ (resp. $\tilde{\theta}_{e^*}$) for all edges incident to this vertex (resp. face), a total angle easily seen to be a multiple of 2π .

Definition 12. An isoradial immersion is said to be *flat* if all the cone angles are equal to 2π , *i.e.*, if for any inner vertex v and any inner face f of \mathbf{G} , we have

$$\sum_{e \sim v} \tilde{\theta}_e = 2\pi \quad \text{and} \quad \sum_{e^* \sim f} \tilde{\theta}_{e^*} = 2\pi. \quad (2)$$

Here are some remarks on the above definitions.

Remark 13.

1. Given the embedding of \mathbf{G}^\diamond in the plane, an isoradial immersion is equivalent to the data of an element $\alpha \in Z_{\mathbf{G}}$ (or equivalently of a rhombic immersion) and of an element $k \in \mathbb{Z}^E$.
2. Let us recall, see for example [Ken02, KS05], that the terminology *isoradial* stems from the fact that in a rhombic immersion of \mathbf{G}^\diamond , the boundary vertices of each primal/dual face corresponding to a dual/primal vertex f/v are mapped to a unit circle whose center is the immersed vertex f/v .
3. The notion of flat isoradial immersion can be formulated very naturally in terms of the graph of train-tracks $\mathbf{G}^{\mathcal{T}}$: the angles α are associated to the train-tracks, the integers k are associated to the vertices of $\mathbf{G}^{\mathcal{T}}$, and they need to verify one condition for each inner vertex v and inner face f of \mathbf{G} , *i.e.*, for each bounded face of $\mathbf{G}^{\mathcal{T}}$:

$$\sum_{e \sim v} k_e = 1 - \frac{1}{2\pi} s_v \quad \text{and} \quad - \sum_{e^* \sim f} k_e = 1 - \frac{1}{2\pi} s_f, \quad (3)$$

where $s_v = \sum_{e \sim v} \theta_e$ and $s_f = \sum_{e^* \sim f} (\pi - \theta_e)$.

4. In the absence of unbounded faces, the notion of flat isoradial immersion is self-dual up to a sign change of k : by exchanging simultaneously the roles on the one

hand of vertices and faces of G , and on the other hand of primal and dual edges in (3), one sees that α and k define an isoradial immersion of G if and only if α and $-k$ define an isoradial immersion of the dual G^* , see also Section 3.4. In the general case, this notion remains locally self-dual away from the unbounded faces.

In the specific case where $k \equiv 0$, *i.e.*, when lifted rhombus angles $\tilde{\theta}_e \in [0, 2\pi)$, and when the latter are moreover restricted to being in $(0, \pi)$, a flat isoradial immersion is an *isoradial embedding* as introduced in [Ken02, KS05]. As mentioned earlier, Kenyon and Schlenker [KS05] prove that a planar embedded graph G admits an isoradial embedding if and only if the train-tracks make neither closed loops nor self-intersections and two train-tracks never intersect more than once. Furthermore, they identify the space of all isoradial embeddings of such a graph G as a subspace of Z_G . The aim of this section and the next is to prove a theorem of the same flavor in the more general setting of isoradial immersions.

3.2 Statement of the main result

The main question answered in this section is that of characterizing the planar embedded graphs admitting a flat isoradial immersion, and of understanding the space of such immersions. Stating this result is the subject of the present section; the proof is provided in Section 3.3.

We need one preliminary remark. Consider a planar embedded graph G , together with angles $\alpha \in Z_G$ and integers $k \in \mathbb{Z}^E$ such that the corresponding isoradial immersion is flat. Moving along the entirety of a fixed train-track $t \in \mathcal{T}$, we encounter vertices of $G^{\mathcal{T}}$, *i.e.*, edges of G : alternatively add $+1$ and -1 to the corresponding integers k_e . Note that vertices corresponding to self-intersections of t receive a contribution of $+1 - 1 = 0$; this can be checked by recalling that train-tracks, also known as *zig-zag paths*, alternately turn right and left when crossing edges. This defines a new set of integers $k' = (k'_e)$, that are said to be obtained from k by a *shift along t* . One easily checks that the isoradial immersion given by the same $\alpha \in Z_G$ and this new set of integers $k' \in \mathbb{Z}^E$ is still flat: indeed, for any fixed face of $G^{\mathcal{T}}$, the vertices of t that belong to this face (if any) can be grouped in pairs of consecutive vertices of t , whose added contributions vanish. (Vertices corresponding to self-intersections will appear twice in this count, with opposite signs). Note that these shift operations are commutative: shifting a set of integers along t_1 and then along t_2 gives the same result as shifting it first along t_2 and then along t_1 .

This motivates the following definition.

Definition 14. Given $\alpha \in Z_G$, two sets of integers $k, k' \in \mathbb{Z}^E$ satisfying the flatness condition (3) are said to be *equivalent* if k' can be obtained from k via a potentially infinite set of shifts along train-tracks, with each train-track supporting a finite number of shifts.

We now have all the ingredients to state our main theorem.

Theorem 15. *An arbitrary planar, embedded graph G admits a flat isoradial immersion if and only if its train-tracks do not form closed loops. When this is the case, for every $\alpha \in Z_G$, there exists $k \in \mathbb{Z}^E$, unique up to equivalence, such that α and k define a flat isoradial immersion of G .*

We mention one immediate reformulation of this theorem. Let us say that two flat isoradial immersions of G are *equivalent* if given by the same α and equivalent k 's. Then, the space of equivalence classes of flat isoradial immersions of G is given by Z_G if its train-tracks do not form closed loops, and is empty otherwise.

Proving Theorem 15 consists in solving the linear system (2). As noted in the introduction, the same system appears in a surprisingly different setting (equations for R-charges for supersymmetric fields). In [Gul08], the author solves this system for a class of \mathbb{Z}^2 -periodic bipartite graphs, containing periodic minimal graphs. Some elements in the proof like, equivalence classes of solutions, evolution of the solution under elementary moves, sketched there without detailed computations, have the same flavor as the arguments we use to tackle the problem in greater generality.

3.3 Proof of Theorem 15

In order to prove Theorem 15, we need two preliminary results. The first, Lemma 16, consists in translating the equations (2) defining flatness, each of which is an equation on a face cycle of G^J , into equations on oriented closed curves of G^J , as defined in Section 2.2. The main content is that, for each closed curve, the flatness equations combine into an equation involving corner vertices only. The second, Proposition 17, computes the evolution of a solution k to the flatness equations (3) under elementary local moves. Although of independent interest, it is used to prove that a graph G having a train-track that forms a closed loop does not have an isoradial immersion. The actual proof of Theorem 15 comes after these two preliminary results.

Flatness equations on closed curves. Recall the definition of closed curves of the graph G^J given at the beginning of Section 2.2, and that of corners. Whereas a closed curve was interpreted there simply as a subgraph because of its description as an \mathbb{F}_2 -chain, we see it here as an oriented curve, bounding a finite collection of oriented faces, possibly with multiplicities, *i.e.*, as a \mathbb{Z} -chain. There are exactly two possible orientations for a closed curve. We choose one of them by orienting arbitrarily one of the edges, and propagating this orientation along the other edges by the drawing procedure presented in Section 2.2.

Following the standard notation coming from algebraic topology, we write this as $c = \partial(\sum_f n_f f)$ with $n_f \in \mathbb{Z}$ vanishing for all but finitely many faces. An example is given in Figure 13.

The following lemma shows that the flatness condition allows to uniquely determine the value of k at a corner of a closed curve as soon as we know its value at all the other corners.




Figure 13: An oriented closed curve c and labelled faces such that $c = \partial(f_1 + f_2 + f_3 + f_4 - f_5)$. Vertices e_1, e_2, e_3, e_4 of the four types are also indicated.

Lemma 16. *Fix angles $\alpha \in Z_G$. Let c be any closed curve in G^J , and e_0 a corner of c . Then, the flatness condition inside all faces appearing with non-zero multiplicity in $c = \partial(\sum_f n_f f)$ allows to express k_{e_0} as an affine function of the values of k at the other corners, which does not depend on the integers associated to vertices that are not corners of c .*

In particular, if the value of k at all corners except e_0 are fixed arbitrarily, then the value at e_0 is determined uniquely.

Proof. Consider the flatness equations (2) for all faces, and add them up with their respective multiplicities n_f . We now show that for all vertices of G^J that are not corners of c , the corresponding variable $\tilde{\theta}_e$ does not appear in the resulting equation. Indeed, there are four types of such vertices, labelled e_1, e_2, e_3, e_4 in Figure 13:

1. vertices disjoint from the faces bounded by c ;
2. vertices of degree 2 of c , where c stays on the same train-track;
3. vertices of degree 4 of c , where c intersects itself transversally;
4. vertices not on c but inside a collection of faces bounded by c .

In the first case, the variable $\tilde{\theta}_e$ does not appear in the equation, so there is nothing to prove. In the second one, using the notation of Figure 14 (left), the four corresponding adjacent faces f_1, f_2, f_3, f_4 satisfy $n_{f_1} = n_{f_2} = n_{f_3} - 1 = n_{f_4} - 1$ and contribute $(n_{f_1} + n_{f_3})\tilde{\theta}_e + (n_{f_2} + n_{f_4})\tilde{\theta}_{e^*} = (n_{f_1} + n_{f_3})\pi$, so $\tilde{\theta}_e$ indeed does not appear. In the third case, the four corresponding adjacent faces f_1, f_2, f_3, f_4 satisfy $n_{f_1} = n_{f_3} = n_{f_2} - 1 = n_{f_4} + 1$ and contribute $(n_{f_1} + n_{f_3})\tilde{\theta}_e + (n_{f_2} + n_{f_4})\tilde{\theta}_{e^*} = (n_{f_1} + n_{f_3})\pi$, as illustrated in Figure 14 (right). Finally, in the last case, the four corresponding adjacent faces all have the same multiplicity n_f and contribute $n_f(2\tilde{\theta}_e + 2\tilde{\theta}_{e^*}) = 2\pi n_f$.

As a consequence, the resulting equation is a linear equation in the integers k_e corresponding to the corners of c . The lemma follows. \square




Figure 14: Contributions of the faces of G^J (i.e., rhombus vertices) around a vertex of a closed oriented curve c .

Elementary local moves. Used in the proof of Theorem 15 and of independent interest, the following proposition describes the evolution of the solution k to flatness under elementary local moves on the graph G^J , which are derived from Reidemeister moves from knot theory, by forgetting information about crossings, see Figure 15.




Figure 15: The three elementary local moves on the graph G^J , analogous to the Reidemeister moves.

Proposition 17. *Let G and G' be two planar, embedded graphs whose train-tracks are related by a sequence of the three local transformations illustrated in Figure 15. Fix $\alpha \in Z_G$ and integers $k \in \mathbb{Z}^E$ defining a flat isoradial immersion of G . Then, there exists $k' \in \mathbb{Z}^E$ such that $\alpha \in Z_G = Z_{G'}$ and k' define a flat isoradial immersion of G' .*




Figure 16: Differing neighborhood in G and G' when passing from one to the other by the three local moves on train-tracks from Figure 15, illustrating the proof of Proposition 17.

Proof. Let G and G' be two planar graphs with train-tracks related by the first elementary move. The notion of flat isoradial immersion being self-dual, it can be assumed that the self-intersection in G^J corresponds to an edge $e = wv \in E$, with v a degree 1 vertex inside the train-track loop (see Figure 16, left). Let us denote by $f \in F$ the face of G

adjacent to v , by $w' \in V'$ the vertex of G' corresponding to $w \in V$, and by $f' \in F'$ the face of G' corresponding to f , as illustrated in Figure 16. Let us first fix $\alpha \in Z_G$ and integers $k \in \mathbb{Z}^E$ defining a flat isoradial immersion of G . Since the two train-track strands intersecting at e belong to the same train-track t , we have $\theta_e = 0$ which by flatness at v implies the value $k_e = 1$. Undoing the train-track loop amounts to removing the corresponding edge e , and rhombus, whose angles are $\tilde{\theta}_e = 2\pi$ and $\tilde{\theta}_{e^*} = -\pi$. This creates an angle defect in G' : with the same angles $\alpha \in Z_G = Z_{G'}$ and integers k restricted to $E' = E \setminus \{e\}$, the angle at w' is equal to $2\pi - 2\pi = 0$, and the angle inside the face f' is given by $2\pi - ((-\pi) + (-\pi)) = 4\pi$. To recover flatness, pick an arbitrary direction on t , and starting at f' , alternatively add $+1$ and -1 to the integers corresponding to the encountered vertices. One easily checks that the resulting $k' \in \mathbb{Z}^{E'}$ defines a flat isoradial immersion of G' . (Note also that choosing the other direction of t gives equivalent integers.) Conversely, let us fix $\alpha \in Z_{G'} = Z_G$ and integers $k' \in \mathbb{Z}^{E'}$ defining a flat isoradial immersion of G' . Flatness at v requires the value $k_e = 1$ for the new rhombus of G , which creates an angle defect of $\pm 2\pi$ in the adjacent vertex w and face f of G . However, flatness can be recovered by the same strategy as above.

Let us now consider two planar graphs G and G' related by the second elementary move. The notion of flat isoradial immersion being self-dual, we can assume that these graphs are locally given as described in the center of Figure 16. We also assume the notation of this figure. Let us fix $\alpha \in Z_G$ and integers $k \in \mathbb{Z}^E$ defining a flat isoradial immersion of G . We claim that the integers $k' \in \mathbb{Z}^{E'}$ given by the restriction of k to $E' = E \setminus \{e_1, e_2\}$ defines a flat isoradial immersion of G' . Indeed, flatness at v implies that the rhombus angles of e_1 and e_2 satisfy $\tilde{\theta}_{e_1} + \tilde{\theta}_{e_2} = 2\pi$; hence, the contribution of these two rhombi to the angle inside the face f_1 is $\tilde{\theta}_{e_1^*} + \tilde{\theta}_{e_2^*} = (\pi - \tilde{\theta}_{e_1}) + (\pi - \tilde{\theta}_{e_2}) = 0$, and similarly for f_2 . Therefore, removing these two rhombi preserves flatness inside the adjacent faces. Furthermore, flatness at v_1 and v_2 amount to equalities $2\pi = \omega_1 + \tilde{\theta}_{e_1}$ and $2\pi = \omega_2 + \tilde{\theta}_{e_2}$ for some $\omega_1, \omega_2 \in \mathbb{R}$ gathering the rest of the angle contributions. The equation $\tilde{\theta}_{e_1} + \tilde{\theta}_{e_2} = 2\pi$ implies $\omega_1 + \omega_2 = 2\pi$, which means flatness at the vertex v' . Conversely, let us fix integers $k' \in \mathbb{Z}^{E'}$ defining a flat isoradial immersion of G' . Flatness at v' means $\omega_1 + \omega_2 = 2\pi$ for some ω_1 (resp. ω_2) gathering angle contributions from the rhombi above (resp. below) the vertex v' . There is a unique integer k_{e_1} (resp. k_{e_2}) such that the corresponding rhombus angle satisfies $\omega_1 + \tilde{\theta}_{e_1} = 2\pi$ (resp. $\omega_2 + \tilde{\theta}_{e_2} = 2\pi$). The corresponding integers $k \in \mathbb{Z}^E$ define an isoradial immersion of G which is flat at v_1 and v_2 by construction; it is also flat at v since the newly defined rhombus angles satisfy $\tilde{\theta}_{e_1} + \tilde{\theta}_{e_2} = 2\pi$. Finally, this equality implies that flatness is preserved at adjacent faces, as above.

Let us finally consider two planar graphs G and G' related by the third elementary move. Without loss of generality, it can be assumed that these graphs are locally given as illustrated in Figure 16, right. Fix $\alpha \in Z_G = Z_{G'}$ and integers $k \in \mathbb{Z}^E$ defining a flat isoradial immersion of G . Using the notation of Figure 16, let $k' \in \mathbb{Z}^{E'}$ be given by $k'_{e'_i} = -k_{e_i}$ for $i = 1, 2, 3$, and $k'_e = k_e$ for the other edges, which are common to G and G' . This choice is motivated by the fact that for any dual edges e and e^* , the




Figure 17: Notation for the third elementary move of Proposition 17.

rhombus angles $\tilde{\theta}_e = \theta_e + 2k_e\pi$ and $\tilde{\theta}_{e^*} = \theta_{e^*} + 2k_{e^*}\pi$ satisfy $\tilde{\theta}_e + \tilde{\theta}_{e^*} = \theta_e + \theta_{e^*} = \pi$, which implies $k_{e^*} = -k_e$. Therefore, we have the equality $\tilde{\theta}_{(e'_i)^*} = \tilde{\theta}_{e_i}$ for $i = 1, 2, 3$. Using this together with the equation $\tilde{\theta}_{e_1} + \tilde{\theta}_{e_2} + \tilde{\theta}_{e_3} = 2\pi$, one easily checks that k' defines a flat isoradial immersion of G' : the corresponding rhombus angle, expressed in terms of $\tilde{\theta}_{e_1}, \tilde{\theta}_{e_2}, \tilde{\theta}_{e_3}$, are illustrated in Figure 17. The converse is proved analogously (or simply using the self-duality of flat isoradial immersions). \square

Proof of Theorem 15. Let us first consider a planar graph G whose train-tracks make at least one closed loop, and assume by means of contradiction that G admits a flat isoradial immersion. Using the three local moves of Figure 15, the corresponding graph of train-tracks can be transformed to contain a subgraph consisting of a simple closed loop crossed by one train-track, as illustrated in Figure 18 (left): fix one closed loop, remove superfluous intersections with this loop using the second and third local moves, and remove self-intersections of the closed loop to make it simple using the first move. Up to reflection along the vertical axis, the corresponding graph G' is locally given as in Figure 18 (right), whose notation we assume. Since G admits a flat isoradial immersion, so does G' by Proposition 17. By flatness at the face f , the rhombus angles satisfy $(\pi - \tilde{\theta}_{e_1}) + (\pi - \tilde{\theta}_{e_2}) = 2\pi$. But this implies the equality $\tilde{\theta}_{e_1} + \tilde{\theta}_{e_2} = 0$, which contradicts the flatness at v .




Figure 18: Proof of Theorem 15: reduction in the case where a train-track of G makes a closed loop. Left: local train-track configuration. Right: corresponding neighborhood in the graph, with v of degree 2 and the other vertex of degree at least 2.

Let us now assume that G is a planar graph whose train-tracks do not form closed loops,

and fix angles $\alpha \in Z_G$. We want to find $k \in \mathbb{Z}^E$ defining a flat isoradial immersion of G . Note that using (the second move of) Proposition 17, it can be assumed that $G^{\mathcal{J}}$ is connected. Take a set of vertices $M \subset V^{\mathcal{J}}$ as in Lemma 4. We now show that, for any choice of integers $(k_e)_{e \in M}$, there exists a unique completion $k = (k_e)_{e \in V^{\mathcal{J}}}$ defining a flat isoradial immersion of G . By Lemma 4, we have a basis $(c_e)_{e \in V^{\mathcal{J}} \setminus M}$ of $\mathcal{E}(G^{\mathcal{J}})$ made of closed curves such that each $e \in V^{\mathcal{J}} \setminus M$ is a corner of c_e and all other corners of c_e belong to M . By Lemma 16, flatness inside the faces determines a unique value for k on the edges corresponding to these remaining vertices. To be more precise, the closed curve c_e endowed with an arbitrary orientation can be written as $c_e = \sum_f n_{ef} \partial f$ with $n_{ef} \in \mathbb{Z}$, and the corresponding \mathbb{Z} -linear combination of the flatness equations for the bounded faces f gives us this result. These oriented closed curves (c_e) form a basis of the (free abelian group of) cycles of $G^{\mathcal{J}}$. Since the boundary of any bounded face f is such a cycle, it can be expressed as $\partial f = \sum_e m_{fe} c_e$ with $m_{fe} \in \mathbb{Z}$. Therefore, the flatness equations for the c_e 's, which we know are satisfied, imply the flatness equations for the faces, *i.e.*, the flatness of the immersion.

Finally, consider a planar graph G whose train-tracks do not form closed loops, and angles $\alpha \in Z_G$ together with $k, k' \in \mathbb{Z}^E$ both defining flat isoradial immersions of G . Fix a set of vertices $M \subset V^{\mathcal{J}}$ as in Lemma 4, and recall the equivalence relation of Definition 14. Since we have an injective map $\tau: M \rightarrow \mathcal{J}$ with $e \in \tau(e)$ for all $e \in M$, there exists $k'' \sim k$ such that k'' and k' coincide on all edges corresponding to elements of M . By Lemma 4, any $e \in V^{\mathcal{J}} \setminus M$ is the corner of some closed curve c_e of $G^{\mathcal{J}}$ such that all corners of c_e but e belong to M . By Lemma 16, flatness implies that k'' and k' coincide on this last corner e as well. Therefore, the family of integers k'' and k' are equal, so k is equivalent to $k' = k''$. \square

3.4 Additional features

This section contains additional features of isoradial immersions. First, restricting to the case where the graph G is finite, we translate results of Theorem 15 into information on the linear system (3) defining flatness; we also give an alternative explicit geometric construction of the solution. Then, we provide a geometric interpretation of folded rhombi of isoradial immersions using the infinite dihedral group. Last, we prove a discrete version of the Gauss–Bonnet formula.

Flatness condition revisited. In this section and the next, we suppose that the graph G is finite and that $G^{\mathcal{J}}$ does not contain closed loops. Recall the linear system (3) defining flatness :

$$\sum_{e \sim v} k_e = 1 - \frac{1}{2\pi} s_v \quad \text{and} \quad - \sum_{e^* \sim f} k_e = 1 - \frac{1}{2\pi} s_f,$$

where $s_v = \sum_{e \sim v} \theta_e$ and $s_f = \sum_{e^* \sim f} (\pi - \theta_e)$. Then, the matrix M corresponding to this linear system has rows indexed by faces of $G^{\mathcal{J}}$, columns indexed by vertices of $G^{\mathcal{J}}$,

or equivalently by edges of \mathbf{G} ; the coefficient $M_{f,e}$ is non zero iff e is on the boundary of the face f ; when this is the case, it is equal to $+1$, resp. -1 , if f corresponds to an inner vertex, resp. an inner face, of \mathbf{G} .

In order to state the translation of Theorem 15 to the linear system, we need one more notation. Given two integer families $k, k' \in \mathbb{Z}^{\mathcal{V}^{\mathcal{J}}}$ related by a shift along a train-track t (recall Definition 14), we let $k^t \in \mathbb{Z}^{\mathcal{E}}$ denote their difference (with an arbitrary sign). Then, rephrasing Theorem 15 in this context yields:

1. The rank of M is equal to $|\mathcal{F}^{\mathcal{J}}|$ (so the associated map is onto).
2. The kernel of M has dimension $|\mathcal{J}| - 1$, and is generated by $\{k^t : t \in \mathcal{J}\}$.

The penultimate paragraph of the proof of Theorem 15, which uses Lemma 4 and Lemma 16, can be seen in this context as a way to explicitly trigonalize the linear system defined by the matrix M .

Geometric construction of solutions. We now propose an alternative geometric construction of a solution to the linear system (3). For every inner vertex v of \mathbf{G} , let δ_v be the function defined on faces of $\mathbf{G}^{\mathcal{J}}$, taking value 1 at the vertex v , and 0 elsewhere; the function δ_f for an inner face f of \mathbf{G} is defined similarly. We explicitly construct functions k^v and k^f such that $Mk^v = \delta_v$ and $Mk^f = \delta_f$, thus proving surjectivity of the map associated to M . Without loss of generality, let us restrict to an inner vertex v of \mathbf{G} since it will be clear from the proof how to proceed for a face f . We use three building blocks.

For every vertex e of $\mathbf{G}^{\mathcal{J}}$ consider the function $\delta_e \in \mathbb{C}^{\mathcal{V}^{\mathcal{J}}}$, taking value 1 at e and 0 elsewhere. Then, if e is not the end point of a loop edge of $\mathbf{G}^{\mathcal{J}}$, it is on the boundary of two faces corresponding to primal vertices and two faces corresponding to dual ones. In this case, $M\delta_e$ is equal to 1 on the two vertices, -1 on the two faces and 0 elsewhere, see Figure 19 (first figure). If the vertex e is the end point of a loop edge, it is a self-intersection point of a single train-track, and the vertex e is on the boundary of two faces corresponding to primal vertices and one corresponding to a dual one (or the reverse). Then $M\delta_e$ takes value 1 at the two primal vertices and -2 at the dual one (or the reverse), see Figure 19 (second figure).

Consider an edge e of $\mathbf{G}^{\mathcal{J}}$ (where the notation e for an edge of $\mathbf{G}^{\mathcal{J}}$ should not be confused with the notation e for a vertex of $\mathbf{G}^{\mathcal{J}}$), then e belongs to a unique train-track t and bounds two faces of $\mathbf{G}^{\mathcal{J}}$ corresponding to a vertex v_e and a face f_e . From the functions (δ_e) , we construct a function k^e such that Mk^e takes value 1 at v_e and -1 at f_e . Let (e_1, e_2) denote the edge e oriented so that v_e is on the left. Consider the path in the train-track t running from e_1 to the boundary without crossing e_2 , and define k^e to be the alternate sum of the functions (δ_e) running over the encountered vertices along the path, starting with a $+1$; note that one could equivalently take a path from e_2 to the boundary of the graph not crossing e_1 . Then, since two consecutive vertices share one face corresponding to a vertex and one to a face of \mathbf{G} (this is because the graph $\mathbf{G}^{\mathcal{J}}$ has




Figure 19: Left: the image $M\delta_e$ around the vertex e of G^J when e is not the end-point of a loop edge (first picture) and when it is (second picture). Right: the image Mk^e when e is a simple edge of G^J (third picture) or a loop edge (fourth picture).

degree 4 vertices), we have that Mk^e is equal to 1 at v_e , -1 at f_e and 0 at all other inner faces of G^J . Note that if e is a loop, all of the above makes sense with minor modifications.

Consider an inner vertex v of G . From the functions (δ_e) we now construct a function k^v such that $Mk^v = \delta_v$. Consider a simple path in the quad-graph G^\diamond from v to a boundary vertex, and the dual edges (which belong to G^J). Define k^v to be the alternate sum of the functions k^e crossing this path, starting from a $+1$. Then, since two consecutive edges bound a common face of G^J corresponding to a primal or a dual vertex of G , we indeed have that $Mk^v = \delta_v$ except on the boundary vertex or face, which does not enter the linear system. A graphical representation of k^v is given in Figure 20 below.

As an immediate consequence, we recover that the rank of M is equal to $|F^J|$.




Figure 20: Graphical representation of the solution k^v : the marked vertices of G^J are those where the function k^v is possibly non-zero.

Remark 18. This concrete way of solving the linear system has a nice geometric interpretation, best explained when the graph G is isoradial, so let us assume that this is the case. Consider angles $\alpha \in Z_G$, then a natural question is to understand how the solution k varies as one of the angles $\alpha(\vec{t}_j)$ moves on the other side of a neighboring angle $\alpha(\vec{t}_i)$. This local transformation is generic in the sense that any pair α, α' of angles in Z_G can be obtained from each other by a sequence of such local moves. This move amounts to exchanging the angles of the directed train-tracks \vec{t}_i, \vec{t}_j crossing at a vertex e in the graph G^J , and this vertex is incident to two faces corresponding to vertices of G and two faces corresponding to faces of G^* . Since the angles $\alpha(\vec{t}_i), \alpha(\vec{t}_j)$ are adjacent, this

amounts to increasing the value $\frac{1}{2\pi}s_v$ by 1 at the two primal vertices and decreasing the value $\frac{1}{2\pi}s_f$ by -1 at the two dual faces, or the opposite. From the first part of the proof involving the function δ_e , we deduce that this local move on the angles has the effect of increasing or decreasing k by 1 at the vertex e .

The next relevant geometric question is: given that the quantity $\frac{1}{2\pi}s_v$ increases or decreases by 1 at a given inner vertex v or face f , how does the solution k vary? What we prove is that the solution gets modified by adding or subtracting 1 along a family of paths of $\mathbf{G}^\mathcal{T}$ to the boundary starting from vertices of a dual path essentially joining the vertex in question to the boundary.

Folded rhombi and the infinite dihedral group. In our definition of isoradial immersions, we have adopted a combinatorial viewpoint. Recall, see Remark 13, that an isoradial immersion is equivalent to the data of an angle map $\alpha \in Z_G$ and of a set of integers $k = (k_e)_e \in \mathbb{Z}^E$. There is a more geometric viewpoint on these choices of lifts, that we now present. As already mentioned, a rhombus with angle $\tilde{\theta}_e = \theta_e \in (0, \pi)$ should be thought of as embedded in the plane: let us denote this *rhombus state* by 1. When the angle $\tilde{\theta}_e$ increases and crosses the value π , the rhombus opens more and more until it folds along the (primal) edge e : let us denote this state by p (for primal). If the angle $\tilde{\theta}_e$ continues to grow and crosses the value 2π , the rhombus shrinks until it folds again, but this time, along the dual edge e^* : let us denote this new state by pd (for primal-dual). Continuing in this way, we see that each angle $\tilde{\theta}_e > 0$ determines a state of the form $pdp \dots$. In the same way, when the angle $\tilde{\theta}_e \in (0, \pi)$ of an embedded rhombus decreases and becomes negative, the rhombus shrinks until it folds along the dual edge e^* , a state denoted by d . Decreasing $\tilde{\theta}_e$ further leads to the rhombus folding again, leading to states of the form $dpd \dots$. Note that folding a rhombus twice along the same (primal or dual) edge does not change its state, a fact denoted by $p^2 = d^2 = 1$.

In a nutshell, the set of states for a given rhombus form the *infinite dihedral group* D_∞ , best understood and presented in the current situation as the free product of two copies of $\mathbb{Z}/2\mathbb{Z}$:

$$D_\infty = \mathbb{Z}/2\mathbb{Z} * \mathbb{Z}/2\mathbb{Z} = \langle p, d \mid p^2 = d^2 = 1 \rangle .$$

The possible deformations of a given rhombus fit nicely in the Cayley graph of this group, with respect to the generators $\{p, d\}$, as illustrated in Figure 21.

There is more structure on the group D_∞ , that sheds an interesting light on isoradial immersions, and on the proof of Theorem 15.

First, there is a natural *orientation homomorphism* $\epsilon: D_\infty \rightarrow \mathbb{Z}/2\mathbb{Z}$ mapping both generators p and d to the non-trivial element of $\mathbb{Z}/2\mathbb{Z}$. A rhombus is *positively oriented* if its state lies in the kernel of ϵ , and *negatively oriented* otherwise. The group $D_\infty^+ = \ker(\epsilon)$ of positively oriented states is isomorphic to \mathbb{Z} (generated by pd or $dp = (pd)^{-1}$), while $D_\infty \setminus D_\infty^+ \simeq pD_\infty^+$ is also in canonical bijection with \mathbb{Z} . This leads to the integer k_e of the combinatorial viewpoint of Section 3.1.

Also, the group D_∞ is endowed with a natural involution determined by $p \mapsto d$ and $d \mapsto p$. It commutes with the orientation homomorphism, and therefore induces an involution




Figure 21: The possible states of a fixed rhombus, as the vertices of the Cayley graph of $D_\infty = \langle p, d \mid p^2 = d^2 = 1 \rangle$, with the corresponding rhombus angles $\tilde{\theta}_e$ and $\tilde{\theta}_{e^*} = \pi - \tilde{\theta}_e$.

on D_∞^+ , which is nothing but $k \mapsto -k$ via the isomorphism $D_\infty^+ \simeq \mathbb{Z}$. This involution corresponds geometrically to exchanging the primal and dual edges of the rhombi. In other words, it is the algebraic counterpart of the geometric duality on the graph G , exhibiting the fact that isoradial immersions are self-dual, *i.e.*, coherent with duality, see Point 4. of Remark 13.

Therefore, instead of defining an isoradial immersion as coming from an angle map $\alpha \in Z_G$ and integers $k \in \mathbb{Z}^E$, we could just as well have defined this object as coming from an $\alpha \in Z_G$ together with some states in $(D_\infty)^E$ with orientations determined by α . As the reader will easily check, the whole of Section 3 can be rewritten in this way, using the duality involution on D_∞ instead of $k \mapsto -k$ (see in particular the third part of the proof of Proposition 17).

Discrete Gauss–Bonnet formula. Consider a planar graph G such that $G^\mathcal{J}$ has no closed loops, and a simple, connected cycle c of $G^\mathcal{J}$. By adding the left-hand side of the flatness equations (2) for all faces inside c , similarly to what we have done in the beginning of the proof of Lemma 16 for oriented closed curves, we obtain a discrete version of the Gauss–Bonnet formula.

Recall the definition of corners given in Section 2.2. In order to state our result, we need the following notation. Denote by $F_c^\mathcal{J}$ the set of faces of $G^\mathcal{J}$ bounded by c . Recall that faces of $G^\mathcal{J}$ correspond to inner vertices and faces of G ; denote by V_c , resp. V_c^* , the set of vertices, resp. faces of G , whose corresponding faces are in $F_c^\mathcal{J}$. Note that each corner vertex e of c is on the boundary of exactly one or three faces of $F_c^\mathcal{J}$; in the first case, we shall write $\text{sgn}(e) = +1$ and in the second case, $\text{sgn}(e) = -1$. Finally, we shall write $e \sim v$ (resp. $e \sim f$) if among the (one or three) faces in $F_c^\mathcal{J}$ bounded by the corner vertex e , exactly one corresponds to a vertex (resp. a face) of G .

Proposition 19 (Discrete Gauss–Bonnet formula). *Fix angles $\alpha \in Z_G$, integers $k = (k_e)_{e \in E}$, and consider the corresponding isoradial immersion. Let c be a simple, connected cycle in $G^\mathcal{J}$. Then, we have the equality*

$$\sum_{v \in V_c} (2\pi - \tilde{s}_v) + \sum_{f \in V_c^*} (2\pi - \tilde{s}_f) + \sum_{e \text{ corner}, e \sim v} \text{sgn}(e) \tilde{\theta}_e + \sum_{e \text{ corner}, e \sim f} \text{sgn}(e) \tilde{\theta}_{e^*} = 2\pi, \quad (4)$$

where $\tilde{s}_v = \sum_{e \sim v} \tilde{\theta}_e$ and $\tilde{s}_f = \sum_{e \sim f} \tilde{\theta}_{e^*}$.

Proof. Throughout this proof, we work in the subgraph $G_c^{\mathcal{J}}$ of $G^{\mathcal{J}}$ bounded by the simple cycle c . There are three types of vertices in this subgraph: vertices of degree 2 in $G_c^{\mathcal{J}}$ (which are corners bounding exactly one face), vertices of degree 3 (which are non-corner boundary vertices), and vertices of degree 4 (which can be either inner vertices, or corners bounding three faces). We denote these set of vertices of $G_c^{\mathcal{J}}$ by V_2, V_3, V_4 respectively.

When summing the left-hand side of equations (2) defining flatness on all faces of $G_c^{\mathcal{J}}$, we have the following contributions:

- every (corner) vertex e of degree 2 is incident to a single face of $G^{\mathcal{J}}$; it contributes $\tilde{\theta}_e$, resp. $\tilde{\theta}_{e^*}$ to the sum if this face corresponds to a vertex, resp. a face of G .
- every vertex e of degree 3 contributes $\tilde{\theta}_e + \tilde{\theta}_{e^*} = \pi$.
- every inner vertex e (of degree 4) contributes $\tilde{\theta}_e + \tilde{\theta}_{e^*} + \tilde{\theta}_e + \tilde{\theta}_{e^*} = 2\pi$.
- every corner vertex e of degree 4 is incident to three faces of $G^{\mathcal{J}}$; it contributes $\tilde{\theta}_{e^*} + \tilde{\theta}_e + \tilde{\theta}_{e^*} = 2\pi - \tilde{\theta}_e$ if only one of these faces corresponds to a vertex of G , and $2\pi - \tilde{\theta}_{e^*}$ in the other case.

As a consequence, we have the equation:

$$\sum_{v \in V_c} \tilde{s}_v + \sum_{f \in V_c^*} \tilde{s}_f = \pi|V_3| + 2\pi|V_4| + \sum_{e \text{ corner}, e \sim v} \text{sgn}(e) \tilde{\theta}_e + \sum_{e \text{ corner}, e \sim f} \text{sgn}(e) \tilde{\theta}_{e^*}.$$

Since $|V_c| + |V_c^*| = |F_c^{\mathcal{J}}|$, we are left with showing that $|V_3| + 2|V_4| = 2|F_c^{\mathcal{J}}| - 2$. This follows from the three equalities $|V_2| + |V_3| + |V_4| = |V_c^{\mathcal{J}}|$, $2|V_2| + 3|V_3| + 4|V_4| = 2|E_c^{\mathcal{J}}|$ and $2|V_2| + |V_3| = 2|\mathcal{J}_c|$ (which already appeared in the proof of Lemma 4), together with the Euler characteristic computation $|V_c^{\mathcal{J}}| - |E_c^{\mathcal{J}}| + |F_c^{\mathcal{J}}| = 1$. This concludes the proof. \square

Remark 20.

- If the isoradial immersion is assumed to be flat, then by definition, we have $\tilde{s}_v = \tilde{s}_f = 2\pi$ for every vertex v and every face f of G . Therefore, Equation (4) simply becomes $\sum_{e \sim v} \text{sgn}(e) \tilde{\theta}_e + \sum_{e \sim f} \text{sgn}(e) \tilde{\theta}_{e^*} = 2\pi$, the sum being over corner vertices.
- If $k \equiv 0$, or equivalently if we consider the rhombic immersion of G^\diamond corresponding to $\alpha \in Z_G$, then the lifted rhombus angles are in $[0, 2\pi)$, and the same equation as (4) holds without the “tilde”. In particular, in the flat case, we get

$$\sum_{e \text{ corner}, e \sim v} \text{sgn}(e) \theta_e + \sum_{e \text{ vertex}, e \sim f} \text{sgn}(e) \theta_{e^*} = 2\pi. \quad (5)$$

It is this special case of our Gauss–Bonnet formula that will prove useful.

4 Minimal immersions of bipartite graphs

As already mentioned, Theorem 15 is in the same spirit as the main result of [KS05] but concerns more general objects, namely flat isoradial immersions of graphs without train-track loops. We now obtain an analogous result for an intermediate class, namely minimal immersions of minimal bipartite graphs, and give an application to the study of the dimer model on such graphs. The main definitions and statements are given in Section 4.1, and the proofs in Section 4.2. Section 4.3 contains our application to dimers.

In all of this section, we suppose that the planar graph G is such that $G^{\vec{\mathcal{J}}}$ does not contain closed loops. For simplicity, we also assume that G does not contain any unbounded faces.

4.1 Minimal immersions

Recall from Section 3.1 that given a planar embedded graph G , an isoradial immersion is built from angles $\alpha \in Z_G$ and integers $k \in \mathbb{Z}^E$; an isoradial immersion is flat if the corresponding total angle at each face and vertex is equal to 2π .

We now assume that G is bipartite (recall Section 2.3). Using the consistent orientation of train-tracks, the space Z_G can be identified with the space of angle maps

$$\alpha: \vec{\mathcal{J}} \rightarrow \mathbb{R}/2\pi\mathbb{Z}.$$

We focus on *flat* isoradial immersions given by such an $\alpha \in Z_G$, but with k identically zero and non-degenerate rhombi. In other words:

Definition 21. A flat isoradial immersion of a planar, embedded bipartite graph G is called a *minimal immersion* of G if all rhombus angles at vertices of G belong to $(0, 2\pi)$.

We give a concrete geometric characterization of minimal immersion in Proposition 25 below, but let us first progress towards stating the main result of this section.

Note that an isoradial immersion with $k = 0$ is fully determined by the angle map $\alpha \in Z_G = \{\vec{\mathcal{J}} \rightarrow \mathbb{R}/2\pi\mathbb{Z}\}$. However, because of the flatness and non-degeneracy conditions, not every α gives a minimal immersion, and not every graph actually admits such an immersion. Our main result characterizes graphs admitting a minimal immersion, and for such a graph G , the subspace of Z_G given by the angles which define such an immersion.

To state this theorem, we need two additional notation. Recall the global cyclic order on $\vec{\mathcal{J}}$ given in Definition 6, and for each vertex v and face f of G , the local cyclic orders on the sets $\vec{\mathcal{J}}(v)$ and $\vec{\mathcal{J}}^\bullet(f), \vec{\mathcal{J}}^\circ(f)$ defined in Section 2.3.

Recall also that a map $\alpha: T \rightarrow \mathbb{R}/2\pi\mathbb{Z}$ on a cyclically ordered set T is called *monotone* if it respects the cyclic order on T and the natural cyclic order on $\mathbb{R}/2\pi\mathbb{Z}$, *i.e.*, whenever we have $[t_1, t_2, t_3]$ in T , then we have $[\alpha(t_1), \alpha(t_2), \alpha(t_3)]$ in $\mathbb{R}/2\pi\mathbb{Z}$. Note that in general we do not ask for strict monotonicity.

Definition 22. Let $X_G \subset Z_G$ denote the set of monotone maps $\alpha: \vec{\mathcal{T}} \rightarrow \mathbb{R}/2\pi\mathbb{Z}$ mapping pairs of intersecting or anti-parallel train-tracks to distinct angles. Similarly, we denote by Y_G the subset of Z_G given by the angle maps whose restriction to $\vec{\mathcal{T}}(v)$ is monotone, injective and non-constant for all vertices v , and whose restrictions to $\vec{\mathcal{T}}^\bullet(f)$ and $\vec{\mathcal{T}}^\circ(f)$ are monotone and non-constant for all faces f of G .

A few remarks are in order. First note that for any graph G , the set X_G is clearly non-empty. Next, as we will prove below, Lemmas 8 and 9 imply that if G is minimal, then we have the inclusion $X_G \subset Y_G$. Since X_G is non-empty, this implies that Y_G is also non-empty for G minimal.

We will actually show that Y_G is non-empty if and only if G is minimal. This is a direct consequence of the main result of this section, that we are now ready to state.

Theorem 23. *An embedded, planar, bipartite graph G admits a minimal immersion if and only if G is minimal. In any case, the space of immersions is given by Y_G .*

The proof of Theorem 23 is postponed to the next section.

Since the space of minimal immersions of a graph is given by Y_G , which contains X_G if G is minimal, it is natural to wonder whether these two spaces coincide for minimal graphs. This turns out to be the case in the \mathbb{Z}^2 -periodic case, see Corollary 29 below.

The following proposition gives the promised geometric characterization of minimal immersions. Prior to stating it, let us make several remarks.

Remark 24.

1. By Point 2. of Remark 13, boundary vertices of primal/dual faces corresponding to dual/primal vertices of G are mapped by the rhombic immersion to points on a circle of unit radius. These points on the unit circle are endowed with a cyclic order, and it makes sense to compare the cyclic order of these vertices in the planar embeddings of G, G^* and the cyclic order of their images in the rhombic immersion of G^\diamond . We use these two orders in Proposition 25 below.
2. If a graph has a vertex of degree 1, then one easily checks that it does not admit any minimal immersion. Therefore, in stating the following geometric description of minimal immersions, we can assume that the graph does not have such a vertex.

Proposition 25. *Let G be a planar, embedded, bipartite graph without degree 1 vertices. Then, a minimal immersion of G is equivalent to a rhombic immersion of G^\diamond satisfying the following conditions:*

1. *For every vertex of G , the boundary vertices of the corresponding face of G^* are all mapped to distinct points, and the cyclic orders on these points coincide in the embedding of G^* and in the rhombic immersion; in particular there is at most one folded rhombus around the vertex.*

2. For every face of \mathbf{G} , the white (resp. black) boundary vertices are not all mapped to the same point, and the cyclic orders on the boundary vertices in the embedding of \mathbf{G} and in the rhombic immersion differ exactly by elementary transpositions given by the boundary edges corresponding to the folded rhombi (in the case of a degree 4 face, there is at most one such edge).

The second point is illustrated in the left side of Figure 22. The proof of this proposition is given in Section 4.2, but let us already state several consequences.

First, folded rhombi in minimal immersions are isolated, in the sense that they meet at most along dual vertices. The right side of Figure 22 shows a typical such folded rhombus, together with adjacent rhombi.

Also, note that the first condition implies that a minimal immersion of \mathbf{G} yields a local embedding of its dual graph \mathbf{G}^* (with the slight abuse that degree 2 faces collapse to a segment) so that each face of \mathbf{G}^* is inscribed in a circle of radius 1. However, this is not an isoradial embedding of \mathbf{G}^* in general, as the circumcenter of a face might not belong to the interior of that face.




Figure 22: Left: a typical minimal immersion around a face. Right: a folded rhombus in a minimal immersion is isolated.

More technically, a minimal immersion of \mathbf{G} defines a flat metric with cone singularities on the surface S given by the union of all (bounded) faces of \mathbf{G} , together with a continuous map $S \rightarrow \mathbb{R}^2$ which is a local isometry away from the “fold lines” given by the boundary of the folded rhombi. The angle at each cone singularity in S can be different from 2π (if it lies on such a fold line), but the average of these angle around any folded rhombus is equal to 2π .

4.2 Proof of Theorem 23, of Proposition 25, and consequences

We begin by identifying the set of maps $\alpha \in Z_{\mathbf{G}}$ defining a flat isoradial immersion of \mathbf{G} .

Lemma 26. *If \mathbf{G} is a bipartite, planar graph, then the isoradial immersion of \mathbf{G} defined by $\alpha \in Z_{\mathbf{G}}$ and $k \equiv 0$ is a minimal immersion if and only if α belongs to $Y_{\mathbf{G}}$.*

Proof. Recall that the isoradial immersion of \mathbf{G} with $k \equiv 0$ is minimal if and only if all the rhombus angles belong to $(0, 2\pi)$, and the sum of the rhombus angles around

every vertex and face of G is equal to 2π . Let us first consider a vertex v of G . If this vertex is of degree 1, then the isoradial immersion of G defined by any $\alpha \in Z_G$ has angle 0 at v , so no angle map α defines a flat immersion. On the other hand, the set $\vec{\mathcal{T}}(v)$ has a single element so the restriction of any $\alpha \in Z_G$ to $\vec{\mathcal{T}}(v)$ is constant; it follows that Y_G is empty. Hence, the lemma holds in this degenerate case, and it can be assumed that v has degree $n \geq 2$. Let us number the adjacent train-tracks strands $\vec{\mathcal{T}}(v)$ and adjacent edges as described in Figure 9 (left). Then, the total angle around v is equal to $\sum_{j=1}^n \theta_{e_j}$ where by definition, the rhombus angle θ_{e_j} is the unique lift in $[0, 2\pi)$ of the angle $[\alpha(\vec{t}_{j+1}) - \alpha(\vec{t}_j)] \in \mathbb{R}/2\pi\mathbb{Z}$. Therefore, the total angle around v is equal to 2π with all rhombus angles in $(0, 2\pi)$ if and only if the restriction of α to $\vec{\mathcal{T}}(v)$ is monotone and injective (hence non-constant since $n > 1$). This is the first condition for α belonging to the set Y_G .

Let us now assume that this condition is satisfied around each vertex v of G (which implies that all the rhombus angles lie in $(0, 2\pi)$), and consider a face f . If this face is of degree 2, then any isoradial immersion of G has a cone angle 0 at f (so the flatness condition is not satisfied) unless the two corresponding train-tracks are assigned identical angles, in which case we have two degenerate rhombi; in any case, this is not a minimal immersion. Also, both sets $\vec{\mathcal{T}}^\bullet(f)$ and $\vec{\mathcal{T}}^\circ(f)$ have a single element, so Y_G is empty. Hence, the statement holds in this case, and it can be assumed that f has degree $2m$ with $m \geq 2$. Numbering the adjacent train-track strands $\vec{\mathcal{T}}(f)$ as in Figure 10, the total angle around the face f is equal to

$$\sum_{j=1}^m (\theta_{e_j^*} + \theta_{e_j'^*}) = \sum_{j=1}^m (2\pi - (\theta_{e_j} + \theta_{e_j'})),$$

where by definition, the angle $\theta_{e_j} \in (0, 2\pi)$ is the lift of $[\alpha(\vec{t}_j) - \alpha(\vec{t}_{j+1})]$ while $\theta_{e_j'} \in (0, 2\pi)$ is the lift of $[\alpha(\vec{t}_j') - \alpha(\vec{t}_{j+1}')]$. Observe that, for every j , the train-tracks $\vec{t}_j, \vec{t}_j', \vec{t}_{j+1}$ are three consecutive train-tracks around a (white) vertex v in the boundary of f . By assumption, the map α restricted to $\vec{\mathcal{T}}(v)$ is monotone, implying that $2\pi - (\theta_{e_j} + \theta_{e_j'})$ is the unique lift in $[0, 2\pi)$ of $[\alpha(\vec{t}_{j+1}') - \alpha(\vec{t}_j)']$. As a consequence, the total angle around f is equal to 2π if and only if the restriction of α to $\vec{\mathcal{T}}^\bullet(f)$ is monotone and non-constant. This is the second condition for α belonging to the set Y_G .

Pairing up the angles around f in the other natural way, and using the first condition around black vertices, we obtain that flatness at f is equivalent to the restriction of α to $\vec{\mathcal{T}}^\circ(f)$ being monotone. This is the last condition for α belonging to the set Y_G , and the proof is completed. \square

Remark 27. It follows from the proof above that if the restriction of $\alpha \in Z_G$ to $\vec{\mathcal{T}}(v)$ is monotone, injective and non-constant for all $v \in V$, then its restriction to $\vec{\mathcal{T}}^\bullet(f)$ is monotone and non-constant for all $f \in F$ if and only if its restriction to $\vec{\mathcal{T}}^\circ(f)$ is monotone and non-constant for all $f \in F$. Therefore, one could drop one of these two conditions in the definition of Y_G .

We are now ready to prove our main result.

Proof of Theorem 23. We first prove the “only if” part using the discrete Gauss–Bonnet formula in the flat case, *i.e.*, Equation (5). Let us assume by means of contradiction that a bipartite planar graph G is not minimal but admits a minimal immersion. Since the graph is not minimal, it either contains a self-intersection or a parallel bigon. In the first case, any self-intersection defines a degenerate rhombus (with angle $\theta = 0$), contradicting the definition of minimal immersion. In the second case, consider the simple closed curve c on G^J defined by an innermost parallel bigon. There are two corners of c , denoted by e_1 and e_2 , each of which is incident to a face of G . By definition of the rhombus angles, we have $\theta_{e_1} = \theta$ and $\theta_{e_2} = 2\pi - \theta$ for some $\theta > 0$. The equality $\theta_{e_1^*} + \theta_{e_2^*} = 0$ follows, contradicting Equation (5). This shows that a non-minimal graph does not admit a minimal immersion. By Lemma 26, this implies that the set Y_G is empty for G non-minimal.

Let us now assume that G is minimal. By Lemma 26, an angle map $\alpha \in Z_G$ defines a minimal immersion of G if and only if α belongs to Y_G , so we are left with proving that Y_G is non-empty. Since X_G is clearly non-empty, it is enough to check the inclusion $X_G \subset Y_G$. Hence, let us fix an element $\alpha \in X_G$, a vertex v of G of degree $n \geq 1$, and study the restriction of α to $\vec{J}(v)$. Since G is minimal, we cannot have $n = 1$. If $n = 2$, then the monotonicity condition is trivially satisfied while the injectivity follows from the fact that the two elements of $\vec{J}(v)$ intersect. For a vertex of degree $n \geq 3$, Lemma 8 together with the definition of X_G readily imply that the restriction of α to $\vec{J}(v)$ is monotone and injective. Let us now turn to a face f of G of degree $2m$, with $m \geq 2$ since G is minimal. Lemma 9 and the definition of X_G immediately imply that the restrictions of α to $\vec{J}^\bullet(f)$ and $\vec{J}^\circ(f)$ are monotone, and non-constant since non-parallel train-tracks either intersect or are anti-parallel. This completes the proof. \square

We now give the proof of the geometric characterization of minimal immersions.

Proof of Proposition 25. Let us fix an embedded bipartite planar graph G without degree 1 vertices. As first requirement, a minimal immersion of G is an isoradial immersion with $k = 0$. By Point 1. of Remark 24, this is equivalent to a rhombic immersion of G^\diamond or to being given an angle map $\alpha \in Z_G$. The second requirement is that it needs to be flat with non-degenerate rhombi, which by Lemma 26 is equivalent to asking that α belongs to Y_G . We are therefore left with checking that Points 1. and 2. are equivalent to $\alpha \in Y_G$.

Using the notation of the proof of Lemma 26, we have that in the rhombic immersion of G^\diamond , boundary vertices of a face of G^* corresponding to a black vertex b of G are immersed as the endpoints of the vectors $e^{i\alpha(\vec{t}_1)}, \dots, e^{i\alpha(\vec{t}_n)}$ drawn in a circle of center b . An analogous statement holds for white vertices with minus those vectors. Translating the first condition of $\alpha \in Y_G$ in the language of the immersed boundary vertices, we get Point 1: the immersion is injective and preserves the cyclic order around this face. As

a consequence of flatness, since a folded rhombus has angle in $(\pi, 2\pi)$, there can be at most one adjacent to any given vertex.

In a similar way, white, resp. black, boundary vertices of a face f of G are immersed as the endpoints of the vectors $e^{i\alpha(\vec{t}_1)}, \dots, e^{i\alpha(\vec{t}_m)}$, resp. $-e^{i\alpha(\vec{t}'_1)}, \dots, -e^{i\alpha(\vec{t}'_m)}$ drawn in a circle of center f . Translating the second condition of $\alpha \in Y_G$ in the language of the immersed boundary vertices exactly says that the immersion preserves the cyclic order of the black, resp. white vertices around this face, which we refer to as Point 2'. We are left with proving that Points 1. and 2. are equivalent to Points 1. and 2'. Around any face f , several folded rhombi can appear, but the corresponding folded edges are disjoint by Point 1. Let us move around the boundary of the immersed face f via the (say, counterclockwise) orientation given by the embedding of G . By construction, an immersed edge in ∂f oriented in this way has the image of the face to its right if and only if it is folded. By flatness, the immersed oriented boundary of f makes exactly one positive turn around the image of the face. One now easily checks that if the cyclic orders on the boundary vertices given by the embedding of G and by the rhombic immersion differ by more than the elementary transpositions given by these folded edges, then Point 1. or Point 2'. is contradicted. (In the case of a face of degree 4, flatness readily implies that there is at most one folded edge on its boundary). Conversely, Point 2. readily implies that the cyclic order of the black, resp. the white vertices is preserved in the immersion, *i.e.* Point 2'., thus ending the proof. \square

As a corollary of Theorem 23, we obtain the following converse to Lemmas 8 and 9.

Corollary 28. *Let G be a bipartite, planar graph without train-track loops and without degree 1 vertices. Let us assume that the conclusions of Lemmas 8 and 9 hold, *i.e.**

1. *For any vertex v of G , all the train-track strands in $\vec{\mathcal{T}}(v)$ belong to distinct train-tracks, and the global cyclic order on $\vec{\mathcal{T}}$ restricts to the local cyclic order on $\vec{\mathcal{T}}(v)$.*
2. *For any face f of G , there is a pair of train-track strands in $\vec{\mathcal{T}}^\bullet(f)$ (resp. $\vec{\mathcal{T}}^\circ(f)$) belonging to distinct non-parallel train-tracks, and the global cyclic order on $\vec{\mathcal{T}}$ restricts to the local cyclic orders on $\vec{\mathcal{T}}^\bullet(f)$ and on $\vec{\mathcal{T}}^\circ(f)$.*

Then, the graph G is minimal.

Proof. Let G be a bipartite planar graph satisfying the assumptions stated above. By the end of the proof of Theorem 23, this implies that the non-empty set X_G is contained in the set Y_G . (Note that G has neither degree 1 vertices, by hypothesis, nor degree 2 faces, by Point 2.) Thus, the set Y_G is non-empty, which by Theorem 23 implies that G is minimal. \square

We conclude this section with the aforementioned result for \mathbb{Z}^2 -periodic minimal graphs. Its proof provides a positive answer to the question at the end of Section 2.3 for this specific class of graphs.

Corollary 29. *If G is minimal and \mathbb{Z}^2 -periodic, then both spaces X_{G} and Y_{G} coincide with the space of monotone maps $\alpha: \vec{\mathcal{T}} \rightarrow \mathbb{R}/2\pi\mathbb{Z}$ mapping intersecting train-tracks to distinct angles.*

Proof. Let us denote by X'_{G} the space of monotone maps $\alpha: \vec{\mathcal{T}} \rightarrow \mathbb{R}/2\pi\mathbb{Z}$ mapping intersecting train-tracks to distinct angles. By definition, we have the inclusion $X_{\mathsf{G}} \subset X'_{\mathsf{G}}$. Furthermore, by the end of the proof of Theorem 23, the inclusion $X_{\mathsf{G}} \subset Y_{\mathsf{G}}$ holds for any minimal graph. Hence, we are left with the proof of the inclusions $Y_{\mathsf{G}} \subset X'_{\mathsf{G}}$ and $X'_{\mathsf{G}} \subset X_{\mathsf{G}}$ for G a \mathbb{Z}^2 -periodic minimal graph. This is an immediate consequence of the following three claims:

1. Any $\alpha \in Y_{\mathsf{G}}$ maps intersecting train-tracks to distinct angles.
2. If α belongs to Y_{G} , then $\alpha: \vec{\mathcal{T}} \rightarrow \mathbb{R}/2\pi\mathbb{Z}$ is monotone.
3. Any $\alpha \in X'_{\mathsf{G}}$ maps (disjoint) anti-parallel train-tracks to distinct angles.

The proof of the first point is straightforward. Indeed, if $\alpha \in Y_{\mathsf{G}}$ were to map intersecting train-tracks to the same angle, then its restriction to $\vec{\mathcal{T}}(v)$ would fail to be injective for v being any of the two vertices bounding the edge corresponding to this intersection.

We now turn to the proof of the second point. To show that $\alpha: \vec{\mathcal{T}} \rightarrow \mathbb{R}/2\pi\mathbb{Z}$ is monotone, the idea is to apply the discrete Gauss-Bonnet formula to the minimal immersion of G given by $\alpha \in Y_{\mathsf{G}}$. Let us fix three pairwise non-parallel train-tracks $\vec{t}, \vec{t}', \vec{t}'' \in \vec{\mathcal{T}}$, and let us first assume that they intersect pairwise. In such a case, there exists an innermost triangle, *i.e.*, a triangle $c \subset \mathsf{G}^{\mathcal{T}}$ whose three sides are segments of the three train-tracks \vec{t}, \vec{t}' and \vec{t}'' , and which does not contain any other such triangle. Since G is minimal, train-tracks do not self-intersect, so this “innermost” condition ensures that c is a simple cycle with three corners of positive signs. Hence, Proposition 19 can be applied in the form of Equation (5): the three rhombus angles corresponding to these three corners add up to 2π . The precise form of these angles depends on the orientation of $\vec{t}, \vec{t}', \vec{t}''$ along c , but the argument is always the same; we shall therefore assume that the three train-tracks $\vec{t}, \vec{t}', \vec{t}''$ are consistently oriented along the curve c (say, counter-clockwise), and leave the other cases to the reader. We shall also assume without loss of generality that c runs along $\vec{t}, \vec{t}', \vec{t}''$ in this (cyclic) order; by minimality of G this implies that these train-tracks are cyclically ordered as $[\vec{t}, \vec{t}', \vec{t}'']$ in $\vec{\mathcal{T}}$. In such a case, Equation (5) reads $2\pi = \theta_e + \theta_{e'} + \theta_{e''}$, where θ_e (resp. $\theta_{e'}$, $\theta_{e''}$) is the unique lift in $(0, 2\pi)$ of the angle $[\alpha(\vec{t}') - \alpha(\vec{t})]$ (resp. $[\alpha(\vec{t}'') - \alpha(\vec{t}')]$, $[\alpha(\vec{t}) - \alpha(\vec{t}'')]$). This equality holds if and only if we have the cyclic order $[\alpha(\vec{t}), \alpha(\vec{t}'), \alpha(\vec{t}'')] in $\mathbb{R}/2\pi\mathbb{Z}$. Therefore, $\alpha: \vec{\mathcal{T}} \rightarrow \mathbb{R}/2\pi\mathbb{Z}$ does preserve the cyclic order in this case.$

Let us now assume that $\vec{t}, \vec{t}', \vec{t}''$ do not intersect pairwise, and finally use the \mathbb{Z}^2 -periodicity of G . In absence of unbounded faces, a train-track is infinite and periodic, so that its quotient by \mathbb{Z}^2 is an oriented closed curve on the torus $\mathbb{R}^2/\mathbb{Z}^2 = \mathbb{T}^2$. We shall write $\Pi(\vec{t}) \in H_1(\mathbb{T}^2; \mathbb{Z}) \simeq \mathbb{Z}^2$ for the homology class of the oriented closed

curve $\vec{t}/\mathbb{Z}^2 \subset \mathbb{T}^2$. Since \mathbf{G} is bipartite and \mathbb{Z}^2 -periodic, the union of all toric, consistently oriented, closed curves coming from train-tracks is the boundary of the faces of $\mathbf{G}^\mathcal{J}/\mathbb{Z}^2 \subset \mathbb{T}^2$ corresponding to vertices of \mathbf{G} ; therefore, the sum of the corresponding homology classes is equal to zero.

Recall that the only way for two closed curves in \mathbb{T}^2 not to meet is if their homology classes coincide (which corresponds to the case of parallel planar curves), or if they are opposite (this gives anti-parallel planar curves). Since $\vec{t}, \vec{t}', \vec{t}''$ are supposed to be pairwise non-parallel and not all to intersect, it can be assumed without loss of generality that \vec{t} and \vec{t}'' are disjoint (with opposite homology classes $\Pi(\vec{t}) + \Pi(\vec{t}'') = 0$), while \vec{t}' meets both \vec{t} and \vec{t}'' . It can also be assumed that these train-tracks are cyclically ordered as $[\vec{t}, \vec{t}', \vec{t}'']$ in $\vec{\mathcal{J}}$; equivalently, the corresponding classes are cyclically ordered as $[\Pi(\vec{t}), \Pi(\vec{t}'), \Pi(\vec{t}'')] in \mathbb{Z}^2 . Since the homology classes of closed curves of \mathbb{T}^2 coming from train-tracks sum to zero, there must exist a train-track homology class $e_1 \in \mathbb{Z}^2$ so that we have the order $[\Pi(\vec{t}), \Pi(\vec{t}'), \Pi(\vec{t}''), e_1]$ in \mathbb{Z}^2 . Let $\vec{t}_1 \in \vec{\mathcal{J}}$ denote a train-track realizing the class e_1 . Translating it far enough in the $\Pi(\vec{t})$ -direction, we can assume that if it intersects \vec{t}' , it does so far away from the (finite number of) intersections of \vec{t}' with $\vec{t} \cup \vec{t}''$. Similarly to the previous case, we can now find a simple closed curve $c \subset \mathbf{G}^\mathcal{J}$ which, when travelled along in the counterclockwise direction, consists of oriented segments of $\vec{t}, \vec{t}_1, \vec{t}'', \vec{t}$, and all of whose (four) corners have positive sign. As in the triangular case above, Equation (5) applied to the curve c readily implies that the cyclic order relation $[\alpha(\vec{t}), \alpha(\vec{t}'), \alpha(\vec{t}''), \alpha(\vec{t}_1)]$ holds in $\mathbb{R}/2\pi\mathbb{Z}$. In particular, we have the order $[\alpha(\vec{t}), \alpha(\vec{t}'), \alpha(\vec{t}'')]$, and the second point is proved.$

We finally turn to the third point, fixing $\alpha \in X_{\mathbf{G}}^l$ and two disjoint anti-parallel train-tracks $\vec{t}, \vec{t}' \in \vec{\mathcal{J}}$. Since \mathbf{G} is bipartite, \mathbb{Z}^2 -periodic and does not have unbounded faces, its train-track homology classes add up to zero and span a 2-dimensional lattice. Hence, there exist train-track homology classes $e_1, e_2 \in \mathbb{Z}^2$, each of which is linearly independent from $\Pi(\vec{t}) = -\Pi(\vec{t}')$, with cyclic order $[\Pi(\vec{t}), e_1, \Pi(\vec{t}'), e_2]$. This implies that any train-tracks $\vec{t}_1, \vec{t}_2 \in \vec{\mathcal{J}}$ realizing these classes will intersect \vec{t} and \vec{t}' and satisfy the cyclic order $[\vec{t}, \vec{t}_1, \vec{t}', \vec{t}_2]$ in $\vec{\mathcal{J}}$. Since α is monotone and maps intersecting train-tracks to distinct angles, we cannot have the equality $\alpha(\vec{t}) = \alpha(\vec{t}')$. This concludes the proof. \square

4.3 The Kasteleyn condition for the dimer model

The goal of this section is to see the implications of the results obtained in Section 4.1 and 4.2 for the dimer model defined on a planar, bipartite graph $\mathbf{G} = (\mathbf{B} \sqcup \mathbf{W}, \mathbf{E})$, see [Ken04] for an overview of the dimer model and references therein. More specifically, we study the implications on the Kasteleyn condition.

Suppose that edges are assigned positive weights $\nu = (\nu_e)_{e \in \mathbf{E}}$. A *dimer configuration* of \mathbf{G} is a *perfect matching*, that is, a subset of edges such that every vertex is incident to exactly one edge of this subset. When the graph \mathbf{G} is finite, the weight of a dimer configuration is the product of the weights of its edges, and the *partition function* is the weighted sum of all the dimer configurations of \mathbf{G} .

In the following, we first give some background on the Kasteleyn condition, then describe and characterize the more general question stemming from our work, and finally give the answer using the aforementioned results of Section 4.1 together with work of Thurston [Thu17].

The Kasteleyn condition. In the study of the dimer model a fundamental tool is the *Kasteleyn matrix*, which is the weighted, oriented, adjacency matrix K associated to (G, ν) , whose rows are indexed by white vertices, columns by black ones, and where the orientation of the edges satisfies *Kasteleyn's condition*, see below. Kasteleyn and Temperley–Fisher's celebrated theorem [Kas61, Kas63, Kas67, TF61] (see also [Per69] for the bipartite setting), states that the absolute value of the determinant of K is equal to the partition function of the dimer model on the weighted graph (G, ν) .

An orientation of the edges of G is equivalent to assigning an element $\omega(e) \in \{-1, +1\}$ to each edge $e = (w, b)$ with the convention that $\omega(e) = +1$ if and only if the orientation of e is, say, from w to b . Then, the coefficient of K corresponding to a white vertex w and a black vertex b is defined by

$$K_{w,b} = \sum_{e=(w,b)} \omega(e) \nu_e,$$

the sum being over all the edges between w and b ; in particular, the coefficient is equal to 0 if there is no such edge. Also, the Kasteleyn condition of [Kas67] can be written as follows. For every bounded face f of G ,

$$\omega(\partial f) := \prod_{e \in \partial f} \omega(e) = -(-1)^{|f|/2}, \quad (6)$$

where $|f|$ denotes the degree of f (which is even since G is bipartite).

An extension of this theory is due to Kuperberg [Kup98]. He proves that the group $\{-1, +1\}$ can be replaced by any subgroup G of \mathbb{C}^* containing $\{-1, +1\}$. The map ω now assigns to each oriented edge $e = (w, b)$ an element $\omega(e) \in G$ so that the same edge with opposite orientation $\bar{e} = (b, w)$ is mapped to $\omega(\bar{e}) = \omega(e)^{-1}$. Kasteleyn's condition on such an ω remains as displayed above, the product being on the set of edges in ∂f oriented consistently (say, counterclockwise around f). The partition function theorem then extends as follows: for any ω satisfying Kasteleyn's condition, the determinant of the associated adjacency matrix K is equal to the dimer partition function, up to multiplication by an element of G . A natural choice of group is $G = S^1$: in this case, the modulus of the determinant of K is equal to the dimer partition function.

In his study of the dimer model on isoradial graphs [Ken02], Kenyon introduces the *discrete $\bar{\partial}$ operator*. The associated matrix is the adjacency matrix K as above, with $\nu_e \in (0, 2)$ given by the length of the dual edge e^* of e , and $\omega(e) \in S^1$ given by the direction of the edge $e = (w, b)$. Using the notation of Section 2.1, the phase $\omega = \omega_\alpha$ can be defined in terms of the train-track angles $\alpha \in Z_G$ as

$$\arg(\omega_\alpha(e)) := C + \arg(e^{i\alpha(\vec{t}_2)} - e^{i\alpha(\vec{t}_1)}) \quad (7)$$

for some constant C . (Note that if $\alpha(\vec{t}_1)$ and $\alpha(\vec{t}_2)$ coincide, then the edge-weight vanishes, which amounts to removing the edge from \mathbf{G} ; alternatively, one can set $\omega_\alpha(e) = -e^{i\alpha(\vec{t}_1)} = -e^{i\alpha(\vec{t}_2)}$). Kenyon shows that the phase ω_α satisfies Kasteleyn's condition (6) for any $\alpha \in Z_{\mathbf{G}}$ defining an isoradial embedding of \mathbf{G} , thus proving that the modulus of the determinant of the discrete $\bar{\partial}$ operator gives the dimer partition function on (\mathbf{G}, ν) .

More general weights have recently been introduced by Fock [Foc15], which exhibit the phase ω_α defined in Equation (7) up to a global sign. The corresponding dimer model in the elliptic case is studied in [BCdT20], where instead of considering isoradial embeddings as in [Ken02], we deal with the more general setting of a bipartite planar graph \mathbf{G} together with an arbitrary angle map $\alpha \in Z_{\mathbf{G}}$. Again we need to know whether the associated phase ω_α satisfies the Kasteleyn condition (6). This, together with the more general framework of this paper, naturally leads to a more general question which is the subject of the next paragraph.

The Kasteleyn condition on rhombic immersions. Let \mathbf{G} be *any* bipartite planar graph and let $\alpha \in Z_{\mathbf{G}}$ be *any* set of angles. Recall that, as far as the present study of the dimer model is concerned, assigning the same angle to train-tracks intersecting at an edge amounts to erasing this edge from \mathbf{G} . For this reason, we can assume without loss of generality that α belongs to the subset $Z'_{\mathbf{G}}$ of $Z_{\mathbf{G}}$ given by maps assigning distinct angles to intersecting train-tracks. Note that, in particular, $Z'_{\mathbf{G}}$ is trivially empty when \mathbf{G} contains a self-intersecting train-track.

A natural question is to describe the subset $K_{\mathbf{G}}$ of $Z'_{\mathbf{G}}$ given by

$$K_{\mathbf{G}} = \{\alpha \in Z'_{\mathbf{G}} \mid \omega_\alpha \text{ satisfies Kasteleyn's condition}\}.$$

In particular, is it possible to characterize the class of graphs \mathbf{G} for which $K_{\mathbf{G}}$ is non-empty?

The following lemma gives a characterization of $K_{\mathbf{G}}$ using the rhombic immersion of \mathbf{G}^\diamond . Note that we will be using the notation of Section 3.1 for the rhombus angles.

Lemma 30. *Let \mathbf{G} be a bipartite, planar graph. Then, for any $\alpha \in Z'_{\mathbf{G}}$, the phase ω_α satisfies the Kasteleyn condition around a face f of \mathbf{G} if and only if, in the rhombic immersion of \mathbf{G}^\diamond defined from α , the angle of the conical singularity at the face f is an odd multiple of 2π , that is, if and only if*

$$\sum_{e^* \sim f} \theta_{e^*} = 2\pi [4\pi].$$

Proof. Fix a planar, bipartite graph \mathbf{G} and $\alpha \in Z'_{\mathbf{G}}$, and consider the corresponding rhombic immersion of \mathbf{G}^\diamond near a given arbitrary bounded face f of \mathbf{G} . Recall that any edge $e \subset \partial f$ defines a rhombus (which can be folded or not), and that all these rhombi are pasted cyclically following the cyclic order of the vertices in ∂f . Recall also that these vertices are on a circle, whose center we denote by $\iota(f)$. Fix an arbitrary black

vertex $b \in \partial f$. Let us write $e' = (w', b)$ and $e'' = (b, w'')$ for the adjacent edges in ∂f oriented counterclockwise around f , write θ', θ'' for the corresponding rhombus angles, and f', f'' for the corresponding faces adjacent to f . By definition, the phase ω_α satisfies

$$\omega_\alpha(e')\omega_\alpha(e'') = \frac{\omega_\alpha(w', b)}{\omega_\alpha(w'', b)} = \frac{e^{i\frac{\pi}{2}} e^{i \arg(\iota(f') - \iota(f))}}{e^{i\frac{\pi}{2}} e^{i \arg(\iota(f) - \iota(f''))}} = -\frac{e^{i \arg(\iota(f') - \iota(f))}}{e^{i \arg(\iota(f'') - \iota(f))}} = -e^{\frac{i}{2}(\theta'^* + \theta''^*)},$$

where $\theta'^* \in (-\pi, \pi]$ denotes the dual rhombus angle of the edge e' at $\iota(f)$, and similarly for θ''^* and e'' . Note that this equality holds whatever the status of the edges e' and e'' (folded or unfolded). Multiplying the equation displayed above for all the $|f|/2$ black vertices of ∂f , we get

$$\omega_\alpha(\partial f) = \prod_{e \in \partial f} \omega_\alpha(e) = (-1)^{|f|/2} e^{\frac{i}{2} \sum_{e^* \sim f} \theta_{e^*}},$$

and the lemma follows. \square

The Kasteleyn condition and minimal graphs. Using Theorem 23 and [Thu17], we now show that K_G is non-empty if and only if G is a minimal graph. The “if” part of this statement extends the aforementioned result of Kenyon [Ken02] which deals with the smaller class of bipartite graphs which can be isoradially embedded in the plane, while the “only if” part provides an alternative proof of the fact that non-minimal graphs do not admit minimal immersions.

Recall that the weights that are of interest to us vanish on an edge if both train-tracks crossing this edge have the same angle. Therefore, we can assume without loss of generality that the train-tracks do not self-intersect.

Theorem 31. *Given a bipartite, planar graph G whose train-tracks do not self-intersect, the set K_G is non-empty if and only if G is minimal. In such a case we have the inclusions $X_G \subset Y_G \subset K_G$.*

While the “if” part of this statement follows readily from Theorem 23, our proof of the other direction relies on two lemmas. The first one deals with two local transformations for bipartite graphs introduced by Kuperberg and called *shrinking/expanding of a 2-valent vertex* and *spider move* in [GK13]. They are illustrated in Figure 23.

Lemma 32. *The space K_G is invariant under shrinking/expanding of 2-valent vertices and spider moves.*

Proof. To prove this result, we will make use of Lemma 30 in the following form: the phase ω_α satisfies the Kasteleyn condition around a face f of G as in Figure 10 if and only if

$$\delta_\alpha(f) := \sum_{j=1}^m (2\pi - (\theta_j + \theta'_j))$$




Figure 23: Shrinking/expanding of a 2-valent (black) vertex, and spider move (with black boundary vertices).




Figure 24: Invariance of K_G under shrinking/expanding of a 2-valent black vertex.

is an odd multiple of 2π , where $\theta_j \in (0, 2\pi)$ is the lift of $[\alpha(\vec{t}'_j) - \alpha(\vec{t}_j)]$ and $\theta'_j \in (0, 2\pi)$ the lift of $[\alpha(\vec{t}_j) - \alpha(\vec{t}'_{j+1})]$ (with $\vec{t}'_{m+1} = \vec{t}'_1$).

Let us start with the invariance of K_G under shrinking/expanding of a 2-valent vertex and, without loss of generality, assume that it is black. Using the notation of Figure 24, we get $\delta_\alpha(f_1) = \delta_\alpha(f'_1) + 2\pi - (\theta_{12} + \theta_{21})$, where $\theta_{12} \in (0, 2\pi)$ is the lift of $[\alpha(\vec{t}_1) - \alpha(\vec{t}_2)]$ and $\theta_{21} \in (0, 2\pi)$ the lift of $[\alpha(\vec{t}_2) - \alpha(\vec{t}_1)]$. This leads to $\theta_{12} + \theta_{21} = 2\pi$, so $\delta_\alpha(f_1)$ and $\delta_\alpha(f'_1)$ coincide. The same argument shows the equality $\delta_\alpha(f_2) = \delta_\alpha(f'_2)$, proving the invariance of K_G under this move.




Figure 25: Invariance of K_G under spider move with black boundary vertices.

We now turn to the invariance of K_G under the spider move, assuming the notation of Figure 25. We handle the case of black boundary vertices, the case of white ones being similar. First note the equality

$$\delta_\alpha(f) = 4\pi - (\theta_{21} + \theta_{14} + \theta_{43} + \theta_{32}) = \delta_\alpha(f'),$$

with $\theta_{ij} \in (0, 2\pi)$ the lift of $[\alpha(\vec{t}_i) - \alpha(\vec{t}_j)]$ for $1 \leq i, j \leq 4$. Therefore, the phase ω_α

satisfies the Kasteleyn condition around f if and only if it does around f' . Moreover, this holds for exactly two of the six cyclic orders of these four train-track angles, namely $[\alpha(\vec{t}_1), \alpha(\vec{t}_2), \alpha(\vec{t}_3), \alpha(\vec{t}_4)]$ and $[\alpha(\vec{t}_1), \alpha(\vec{t}_4), \alpha(\vec{t}_3), \alpha(\vec{t}_2)]$. Therefore, it can be assumed that these four evaluations of α are cyclically ordered in one of these two ways. Let us now study the Kasteleyn condition around the faces f_1 and f'_1 . Using the same notation as above, we have

$$\delta_\alpha(f_1) - \delta_\alpha(f'_1) = \theta_{43} + \theta_{31} - (\theta_{21} + \theta_{32}).$$

As one easily checks, the cyclic order constraint imposed on α by the Kasteleyn condition around f implies the equality $\theta_{43} + \theta_{31} = \theta_{21} + \theta_{32}$. Therefore, $\delta_\alpha(f_1)$ and $\delta_\alpha(f'_1)$ coincide, so the phase ω_α satisfies the Kasteleyn condition around f_1 if and only if it does around f'_1 . The equalities $\delta_\alpha(f_j) = \delta_\alpha(f'_j)$ for $j = 2, 3, 4$ can be checked in the same way, proving the invariance of K_G under this move. \square

We now turn to the second and last lemma needed for the proof of Theorem 31.

Lemma 33. *Let G be a bipartite, planar, non-minimal graph whose train-tracks do not self-intersect. Then, it is related via a finite number of shrinking/expanding of 2-valent vertices and spider moves to a bipartite, planar graph G' which admits a face of degree 2.*

Proof. Since G is non-minimal, connected, and its train-tracks do not self-intersect, it contains a finite subgraph G_1 whose train-tracks form at least one parallel bigon. Since G_1 is finite and has no 1-valent vertex (as its train-tracks do not self-intersect), it can be transformed into a graph G_2 whose white vertices are 3-valent, via a finite number of shrinking of (white) 2-valent vertices and expanding of (black) 2-valent vertices. Following [GK13], consider the *triple-crossing diagram* D_2 obtained from the train-tracks \vec{T} of G_2 as follows: for each white vertex w of G , move the three oriented train-tracks of $\vec{T}(w)$ so that they intersect into a single 6-valent vertex. Fix an innermost parallel bigon B in D_2 , and consider the triple-crossing diagram $D \subset D_2$ given by all the train-tracks strands contained in some small neighborhood of B . The proof of [Thu17, Lemma 12] now implies the following: via a finite number of so-called *2–2 moves*, illustrated in Figure 26 (left), one can transform D so that the parallel bigon B no longer contains crossings in its interior, and so that one of the strands bounding B no longer forms anti-parallel bigons in B . This is illustrated in Figure 26 (right). Note that two




Figure 26: Thurston’s 2–2 move (left), and a parallel bigon with neither interior crossings nor anti-parallel bigons bounded by the upper strand (right).

other moves appear in [Thu17, Lemma 12], but they are only required in order to get rid of self-intersections, parallel bigons and disconnected components. Since D admits neither self-intersections nor loops, and since B is innermost among parallel bigons, these other moves are not required in our situation. By [Thu17, Lemma 11(c)] (see also Figure 3(c) of this same article), a finite number of additional 2–2 moves transforms D into a triple-crossing diagram D' containing a diagram of the form illustrated in Figure 27 (left). Now observe that the two types of 2–2 moves given by the two possible orientations of the train-tracks correspond to shrinking/expanding of (black) 2-valent vertices and spider moves (with black boundary vertices), see Figures 24 and 25. Finally, the graph G' corresponding to D' contains the graph illustrated in Figure 27 (right), which admits a face of degree 2. This concludes the proof. \square




Figure 27: A portion of the triple-crossing diagram D' (left), and the corresponding subgraph of G' (right).

As noticed by the reader, our proof of Lemma 33 relies on Thurston’s theory of triple-crossing diagrams [Thu17]. It should be noted that this result also follows from the work of Postnikov, more precisely [Pos06, Lemma 13.6] applied to connected, bipartite planar graphs whose train-tracks do not self-intersect. We thank Marianna Russkikh for pointing out this additional reference.

We are finally ready to prove the main result of this section.

Proof of Theorem 31. Let us first assume that G is minimal. By Theorem 23, any $\alpha \in Y_G$ defines a flat isoradial immersion of G . In particular, the angle of the conical singularity at any face is equal to 2π . By Lemma 30, this shows that ω_α satisfies the Kasteleyn condition. Since the inclusion $Y_G \subset Z'_G$ is immediate, we have the inclusion $Y_G \subset K_G$. By the proof of Theorem 23, we also have the inclusion $X_G \subset Y_G$, with X_G non-empty, so K_G is non-empty as claimed. Let us now assume that G is a non-minimal graph whose train-tracks do not self-intersect. By Lemmas 32 and 33, the space K_G coincides with $K_{G'}$ for some planar bipartite graph G' admitting a face of degree 2. Let us denote by \vec{t}_1, \vec{t}_2 the two corresponding train-tracks. For any $\alpha \in Z'_G = Z'_{G'}$, the cone angle inside this face is equal to $2\pi - (\theta + \theta')$, where $\theta \in (0, 2\pi)$ is the lift of $[\alpha(\vec{t}_1) - \alpha(\vec{t}_2)]$ and $\theta' \in (0, 2\pi)$ the lift of $[\alpha(\vec{t}_2) - \alpha(\vec{t}_1)]$. This leads to $\theta + \theta' = 2\pi$, so this cone angle

vanishes. By Lemma 30, this arbitrary α does not belong to $K_{G'}$, so $K_{G'} = K_G$ is empty. This concludes the proof. \square

A natural question is whether the inclusions $X_G \subset Y_G \subset K_G$ for G minimal are actually equalities. The equality between X_G and Y_G holds in the \mathbb{Z}^2 -periodic case (recall Corollary 29), but whether it remains true in the general case is still open. We conclude this article with an example showing that the second inclusion is not an equality in general.

Example 34. Consider the bipartite minimal graph G given by the square lattice, and denote the four train-tracks around an arbitrary face by t_1, t'_1, t_2, t'_2 , as illustrated in Figure 10. We focus on \mathbb{Z}^2 -periodic elements of Z_G , *i.e.*, maps assigning the same value to parallel train-tracks. Since each train-track of G is parallel to exactly one of the four train-tracks listed above, we are dealing with maps $\alpha: \{t_1, t'_1, t_2, t'_2\} \rightarrow \mathbb{R}/2\pi\mathbb{Z}$. It is not difficult to check that, on this easy example, the set of \mathbb{Z}^2 -periodic elements of Y_G consists of all injective maps α such that the angles are cyclically ordered as $[\alpha(\vec{t}_1), \alpha(\vec{t}'_1), \alpha(\vec{t}_2), \alpha(\vec{t}'_2)]$. However, any injective $\beta: \{t_1, t'_1, t_2, t'_2\} \rightarrow \mathbb{R}/2\pi\mathbb{Z}$ such that the corresponding angles are cyclically ordered as $[\beta(\vec{t}'_1), \beta(\vec{t}_1), \beta(\vec{t}'_2), \beta(\vec{t}_2)]$ will define an isoradial immersion with cone angle equal to -2π inside each face of G . Therefore, by Lemma 30, such a β belongs to $K_G \setminus Y_G$.

References

- [Bax78] Rodney J. Baxter. Solvable eight-vertex model on an arbitrary planar lattice. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 289(1359):315–346, 1978.
- [Bax89] Rodney J. Baxter. *Exactly solved models in statistical mechanics*. Academic Press Inc. [Harcourt Brace Jovanovich Publishers], London, 1989. Reprint of the 1982 original.
- [BCdT20] Cédric Boutillier, David Cimasoni, and Béatrice de Tilière. Minimal bipartite dimers and genus one Harnack curves. *arXiv e-prints*, 2020.
- [BCdT21] Cédric Boutillier, David Cimasoni, and Béatrice de Tilière. Dimers on minimal graphs and genus g Harnack curves, 2021+.
- [BDCS15] Vincent Beffara, Hugo Duminil-Copin, and Stanislav Smirnov. On the critical parameters of the $q \leq 4$ random-cluster model on isoradial graphs. *J. Phys. A*, 48(48):484003, 25, 2015.
- [BdT10] Cédric Boutillier and Béatrice de Tilière. The critical \mathbf{Z} -invariant Ising model via dimers: the periodic case. *Probab. Theory Related Fields*, 147(3-4):379–413, 2010.

- [BdT11] Cédric Boutillier and Béatrice de Tilière. The critical Z -invariant Ising model via dimers: locality property. *Comm. Math. Phys.*, 301(2):473–516, 2011.
- [BdT12] Cédric Boutillier and Béatrice de Tilière. Statistical mechanics on isoradial graphs. In *Probability in complex physical systems*, volume 11 of *Springer Proc. Math.*, pages 491–512. Springer, Heidelberg, 2012.
- [BdTR17] Cédric Boutillier, Béatrice de Tilière, and Kilian Raschel. The Z -invariant massive Laplacian on isoradial graphs. *Invent. Math.*, 208(1):109–189, 2017.
- [BdTR19] Cédric Boutillier, Béatrice de Tilière, and Kilian Raschel. The Z -invariant Ising model via dimers. *Probab. Theory Related Fields*, 174(1-2):235–305, 2019.
- [Cim12] David Cimasoni. The critical Ising model via Kac-Ward matrices. *Comm. Math. Phys.*, 316(1):99–126, 2012.
- [CS11] Dmitry Chelkak and Stanislav Smirnov. Discrete complex analysis on isoradial graphs. *Adv. Math.*, 228(3):1590–1630, 2011.
- [CS12] Dmitry Chelkak and Stanislav Smirnov. Universality in the 2D Ising model and conformal invariance of fermionic observables. *Invent. Math.*, 189(3):515–580, 2012.
- [DCLM18] Hugo Duminil-Copin, Jhih-Huang Li, and Ioan Manolescu. Universality for the random-cluster model on isoradial graphs. *Electron. J. Probab.*, 23:Paper No. 96, 70, 2018.
- [dT07] Béatrice de Tilière. Scaling limit of isoradial dimer models and the case of triangular quadri-tilings. *Ann. Inst. H. Poincaré Sect. B.*, 43(6):729–750, 2007.
- [Duf68] Richard J. Duffin. Potential theory on a rhombic lattice. *J. Combinatorial Theory*, 5:258–272, 1968.
- [FHV⁺06] Sebastián Franco, Amihay Hanany, David Vegh, Brian Wecht, and Kristian D. Kennaway. Brane dimers and quiver gauge theories. *Journal of High Energy Physics*, 2006(01):096–096, jan 2006.
- [Foc15] Vladimir V. Fock. Inverse spectral problem for GK integrable system. *ArXiv e-prints*, March 2015.
- [GK13] Alexander B. Goncharov and Richard Kenyon. Dimers and cluster integrable systems. *Ann. Sci. Éc. Norm. Supér.*, 46(5):747–813, 2013.
- [GM14] Geoffrey R. Grimmett and Ioan Manolescu. Bond percolation on isoradial graphs: criticality and universality. *Probab. Theory Related Fields*, 159(1-2):273–327, 2014.

- [Gri14] Geoffrey R. Grimmett. Criticality, universality, and isoradiality. In *Proceedings of the International Congress of Mathematicians—Seoul 2014. Vol. IV*, pages 25–47. Kyung Moon Sa, Seoul, 2014.
- [Gul08] Daniel R. Gulotta. Properly ordered dimers, R-charges, and an efficient inverse algorithm. *Journal of High Energy Physics*, 2008(10):014, 2008.
- [IU11] Akira Ishii and Kazushi Ueda. A note on consistency conditions on dimer models. *Higher Dimensional Algebraic Geometry*, pages 143–164, 2011.
- [Kas61] Pieter W. Kasteleyn. The statistics of dimers on a lattice. *Physica*, 27:1209–1225, 1961.
- [Kas63] Pieter W. Kasteleyn. Dimer statistics and phase transitions. *J. Math. Phys.*, 4:287–293, 1963.
- [Kas67] Pieter W. Kasteleyn. Graph theory and crystal physics. In *Graph Theory and Theoretical Physics*, pages 43–110. Academic Press, London, 1967.
- [Ken02] Richard Kenyon. The Laplacian and Dirac operators on critical planar graphs. *Invent. Math.*, 150(2):409–439, 2002.
- [Ken04] Richard Kenyon. An introduction to the dimer model. In *School and Conference on Probability Theory*, ICTP Lect. Notes, XVII, pages 267–304. Abdus Salam Int. Cent. Theoret. Phys., Trieste, 2004.
- [KLRR18] Richard Kenyon, Wai Yeung Lam, Sanjay Ramassamy, and Marianna Russkikh. Dimers and Circle patterns. *arXiv e-prints*, page arXiv:1810.05616v1, October 2018.
- [KS05] Richard Kenyon and Jean-Marc Schlenker. Rhombic embeddings of planar quad-graphs. *Trans. Amer. Math. Soc.*, 357(9):3443–3458 (electronic), 2005.
- [Kup98] Greg Kuperberg. An exploration of the permanent-determinant method. *Electron. J. Combin.*, 5:Research Paper 46, 34, 1998.
- [Li17] Zhongyang Li. Conformal invariance of dimer heights on isoradial double graphs. *Ann. Inst. Henri Poincaré D*, 4(3):273–307, 2017.
- [Mer01] Christian Mercat. Discrete Riemann surfaces and the Ising model. *Comm. Math. Phys.*, 218(1):177–216, 2001.
- [Per69] Jerome K. Percus. One more technique for the dimer problem. *J. Math. Phys.*, 10:1881, 1969.
- [Pos06] Alexander Postnikov. Total positivity, Grassmannians, and networks. *arXiv e-prints*, September 2006.

- [TF61] Harold N. V. Temperley and Michael E. Fisher. Dimer problem in statistical mechanics-an exact result. *Philosophical Magazine*, 6(68):1061–1063, 1961.
- [Thu17] Dylan P. Thurston. From dominoes to hexagons. In *Proceedings of the 2014 Maui and 2015 Qinhuangdao conferences in honour of Vaughan F. R. Jones' 60th birthday*, volume 46 of *Proc. Centre Math. Appl. Austral. Nat. Univ.*, pages 399–414. Austral. Nat. Univ., Canberra, 2017.